Third Edition

GUIDE TO NETWORK DEFENSE AND COUNTERMEASURES

Randy Weaver, Dawn Weaver, Dean Farwood

PREPARING TOMORROW'S INFORMATION SECURITY PROFESSIONALS

This is an electronic version of the print textbook. Due to electronic rights restrictions, some third party content may be suppressed. Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. The publisher reserves the right to remove content from this title at any time if subsequent rights restrictions require it. For valuable information on pricing, previous editions, changes to current editions, and alternate formats, please visit www.cengage.com/highered to search by ISBN#, author, title, or keyword for materials in your areas of interest.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Guide to Network Defense and Countermeasures
Third Edition

Randy Weaver

Dawn Weaver

Dean Farwood

COURSE TECHNOLOGY
CENGAGE Learning®

Australia • Brazil • Japan • Korea • Mexico • Singapore • Spain • United Kingdom • United States

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

COURSE TECHNOLOGY
CENGAGE Learning

Guide to Network Defense and Countermeasures, Third Edition

Randy Weaver, Dawn Weaver, Dean Farwood

Vice President, Careers & Computing: Dave Garza

Executive Editor: Steve Helba

Director, Development – Careers & Computing: Marah Bellegarde

Product Development Manager: Juliet Steiner

Product Manager: Natalie Pashoukos

Developmental Editor: Dan Seiter

Editorial Assistant: Jennifer Wheaton

Vice President, Marketing: Jennifer Ann Baker

Marketing Director: Deborah Yarnell

Production Director: Wendy A. Troeger

Production Manager: Andrew Crouth

Senior Content Project Manager: Andrea Majot

Senior Art Director: Jack Pendleton

Technology Project Manager: Joe Pliss

Media Editor: William Overocker

© 2014 Course Technology, Cengage Learning

ALL RIGHTS RESERVED. No part of this work covered by the copyright herein may be reproduced, transmitted, stored or used in any form or by any means graphic, electronic, or mechanical, including but not limited to photocopying, recording, scanning, digitizing, taping, Web distribution, information networks, or information storage and retrieval systems, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the publisher.

For product information and technology assistance, contact us at
Cengage Learning Customer & Sales Support, 1-800-354-9706

For permission to use material from this text or product, submit all requests online at **cengage.com/permissions**
Further permissions questions can be emailed to
permissionrequest@cengage.com

Library of Congress Control Number: 2012949396

ISBN-13: 978-1-133-72794-1

ISBN-10: 1-133-72794-8

Course Technology
20 Channel Center Street
Boston, MA 02210
USA

Cengage Learning is a leading provider of customized learning solutions with office locations around the globe, including Singapore, the United Kingdom, Australia, Mexico, Brazil, and Japan. Locate your local office at: **international.cengage.com/region**

Cengage Learning products are represented in Canada by Nelson Education, Ltd.

For your lifelong learning solutions, visit **www.cengage.com/coursetechnology**

Purchase any of our products at your local college store or at our preferred online store **www.cengagebrain.com**

Visit our corporate website at **cengage.com**.

Microsoft and the Office logo are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Course Technology, a part of Cengage Learning, is an independent entity from the Microsoft Corporation, and not affiliated with Microsoft in any manner.

Course Technology, a part of Cengage Learning, reserves the right to revise this publication and make changes from time to time in its content without notice.

The programs in this book are for instructional purposes only. They have been tested with care, but are not guaranteed for any particular intent beyond educational purposes. The author and the publisher do not offer any warranties or representations, nor do they accept any liabilities with respect to the programs.

Printed in the United States of America
1 2 3 4 5 6 7 16 15 14 13 12

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Brief Contents

INTRODUCTION xv

CHAPTER 1
Network Security Fundamentals 1

CHAPTER 2
TCP/IP 35

CHAPTER 3
Network Traffic Signatures 79

CHAPTER 4
Routing Fundamentals 119

CHAPTER 5
Cryptography 157

CHAPTER 6
Wireless Network Fundamentals 193

CHAPTER 7
Understanding Wireless Security 227

CHAPTER 8
Intrusion Detection and Prevention Systems 265

CHAPTER 9
Firewalls 305

CHAPTER 10
Firewall Design and Management 343

CHAPTER 11
VPN Concepts 385

CHAPTER 12
Internet and World Wide Web Security 437

CHAPTER 13
Security Policy Design and Implementation 475

CHAPTER 14
Ongoing Security Management 525

APPENDIX A 549

GLOSSARY 553

INDEX 567

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Table of Contents

INTRODUCTION xv

CHAPTER 1
Network Security Fundamentals 1

Examining Network Security Fundamentals 2
Threats to Network Security 2
Common Attacks and Defenses 6
Goals of Network Security 9

Using a Layered Defense Strategy: Defense in Depth 11
Physical Security 12
Authentication and Password Security 12
Operating System Security 13
Antivirus Protection 13
Packet Filtering 13
Firewalls 14
Demilitarized Zone (DMZ) 15
Intrusion Detection and Prevention System (IDPS) 15
Virtual Private Networks (VPNs) 16
Network Auditing and Log Files 17
Routing and Access Control Methods 18

The Impact of Defense 19

Chapter Summary 20

Key Terms 20

Review Questions 22

Hands-On Projects 25

Case Projects 33

CHAPTER 2
TCP/IP 35

The OSI Model and TCP/IP Protocols 36
The OSI Model 36
TCP/IP Addressing 37
Address Classes 38
Private IP Address Ranges 38
Subnetting 39
Variable Length Subnet Masking 42
Classless Interdomain Routing 42
Unicasting, Multicasting, and Broadcasting 43

Examining Internet Protocol Version 4 (IPv4) 43
IP Datagrams 43
IP Header Structure 44
ICMP Messages 46
TCP Headers 47
UDP Headers 48
Packet Fragmentation 49
The TCP Life Cycle and the TCP Three-Way Handshake 51
Domain Name System 53

Internet Protocol Version 6 (IPv6) 54
IPv6 Core Protocols 55
IPv6 Addressing 62

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

IPv6 Configuration 63
IPv6 Utilities 63

Chapter Summary 65

Key Terms 67

Review Questions 68

Hands-On Projects 71

Case Projects 77

CHAPTER 3
Network Traffic Signatures 79

Examining the Common Vulnerabilities and Exposures Standard 80
How CVE Works 80
Scanning CVE Vulnerability Descriptions 82

Understanding Signature Analysis 83
Bad Header Information 83
Suspicious Data Payload 84
Single-Packet Attacks 85
Multiple-Packet Attacks 85
Analyzing Packets 85

Analyzing Traffic Signatures 88
Examining Normal Network Traffic Signatures 89
Examining Abnormal Network Traffic Signatures 95

Identifying Suspicious Events 100
Packet Header Discrepancies 101
Advanced Attacks 104
Remote Procedure Call Attacks 105

Chapter Summary 105

Key Terms 106

Review Questions 107

Hands-On Projects 110

Case Projects 118

CHAPTER 4
Routing Fundamentals 119

Examining the Routing Process 120
The Address Resolution Protocol Processes 121
Accessing a Router 121
Routing Tables 122
Static Routing 122
Dynamic Routing 124
Routing Metrics 124
Choosing a Routing Protocol 125
Route Summarization 125

Router Security Fundamentals 128
Creating and Using Access Control Lists 128
Use and Rules 129
Standard ACLs 130
Extended ACLs 131
Named ACLs 132
Examining Cisco Router Logging 133

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Cisco Authentication and Authorization 136
Router Passwords 137
Banners 139
Remote Access with Secure Shell 140
Hardening a Router 143

Chapter Summary 144

Key Terms 145

Review Questions 147

Hands-On Projects 149

Case Projects 156

CHAPTER 5
Cryptography 157

Components of Cryptographic Protocols 158
Cryptographic Primitives 158
Encryption Algorithms 162
Hashing Algorithms 165
Message Authentication Code 167
Digital Signatures 168
Key Management 169

Examining Cryptography Standards 173
Data Encryption Standard 173
Triple DES 173
Advanced Encryption Standard 174
Internet and Web Standards 174
Internet Protocol Security 175

Modern Cryptanalysis Methods 179
Side Channel Attacks 179
Passive Attacks 180
Chosen Ciphertext and Chosen Plaintext Attacks 180
XSL Attacks 181
Random Number Generator Attacks 181
Related Key Attacks 181
Integral Cryptanalysis 182
Differential Cryptanalysis 182

Chapter Summary 183

Key Terms 184

Review Questions 186

Hands-On Projects 188

Case Projects 192

CHAPTER 6
Wireless Network Fundamentals 193

Wireless Communications Primer 194
Electromagnetic Radiation 194
Infrared Transmissions 195
Radio Frequency Transmissions 196

Wireless LANs and Their Components 206
Wireless NICs 206
Access Points 206
Antennas 208

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Remote Wireless Bridges . . . 210
Wireless Gateways . . . 211
WLAN Configurations . . . 212

Wireless Networking Standards . . . 213
IEEE 802.11 . . . 213
Radio Frequency and the FCC . . . 215

Chapter Summary . . . 215

Key Terms . . . 217

Review Questions . . . 220

Hands-On Projects . . . 222

Case Projects . . . 226

CHAPTER 7
Understanding Wireless Security . . . 227

Security Concerns of Wireless Networking . . . 228
IEEE 802.11 Media Access Control: Frames . . . 228
Scanning and Attacks . . . 232
Wardriving and Exploitation of Rogue Devices . . . 234
Wireless Man-in-the-Middle Attacks . . . 235

Secure WLAN Implementation . . . 235
Association with a Wireless Network . . . 235
Wireless Authentication . . . 236
Default WEP Keys . . . 239
Key Management Concerns in 802.11 Networks . . . 240
MAC Address Filtering and Spoofing . . . 240
Wireless Device Portability . . . 240

Examining Wireless Security Solutions and Countermeasures . . . 241
Incorporating a Wireless Security Policy . . . 241
Ensuring Physical Security . . . 242
Planning AP Placement . . . 242
Changing Default Hardware and Software Settings . . . 243
Strong Encryption and Authentication . . . 244
Wireless Auditing . . . 249
AP Logging Functions . . . 250
Best Practices for Wireless Network Security . . . 251
Mobile Device Security . . . 252
Approaches to Mobile Device Security . . . 253

Chapter Summary . . . 254

Key Terms . . . 255

Review Questions . . . 258

Hands-On Projects . . . 260

Case Projects . . . 262

CHAPTER 8
Intrusion Detection and Prevention Systems . . . 265

Goals of an IDPS . . . 266

Common Detection Methodologies . . . 267
Anomaly and Signature Detection Systems . . . 267
Stateful Protocol Analysis . . . 269

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Examining IDPS Components 270
Sensors and Agents 270
Detection and Prevention Capabilities 272
Command Console 273
Database of Attack Signatures or Behaviors 273

Options for IDPSs 274
Network-Based IDPSs 275
Host-Based IDPSs 279
Comparing an NIDPS and HIDPS 282
Hybrid IDPSs 282

Securing IDPS Components 283
IDPS Security Best Practices 284

Developing IDPS Filter Rules 284

Examining Intrusion Detection Step by Step 285
Step 1: Installing the IDPS Database 285
Step 2: Gathering Data 286
Step 3: Sending Alert Messages 287
Step 4: The IDPS Responds 287
Step 5: The Administrator Assesses Damage 287
Step 6: Following Escalation Procedures 288
Step 7: Logging and Reviewing Events 289
Evaluating IDPS Products 289

Chapter Summary 290

Key Terms 290

Review Questions 291

Hands-On Projects 294

Case Projects 303

CHAPTER 9
Firewalls 305

Overview of Firewalls 306
Comparing Software and Hardware Firewalls 308
Software-Based Firewalls 309
Hardware Firewalls 310

Packet Filtering and Firewall Rule Sets 311
Stateless Packet Filtering 312
Stateful Packet Filtering 313
Packet Filtering Based on Position 314
Firewall Rule Sets 317

Chapter Summary 329

Key Terms 330

Review Questions 330

Hands-On Projects 332

Case Projects 340

CHAPTER 10
Firewall Design and Management 343

Designing Firewall Configurations 344
Screening Routers 344
Dual-Homed Hosts 344

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Screened Hosts . . . 346
Screened Subnet DMZs . . . 347
Multiple DMZ/Firewall Configurations . . . 348
Multiple Firewall Configurations . . . 350
Reverse Firewalls . . . 351
Choosing a Firewall Configuration . . . 353

Examining Proxy Servers . . . 353
Goals of Proxy Servers . . . 354
How Proxy Servers Work . . . 355
Choosing a Proxy Server . . . 356
Filtering Content . . . 358

Choosing a Bastion Host . . . 358
General Requirements . . . 359
Selecting the Bastion Host Machine . . . 359
Choosing an Operating System . . . 359
Memory and Processor Speed . . . 360
Location on the Network . . . 360
Hardening the Bastion Host . . . 360
Selecting Services to Provide . . . 361
Using Honeypots . . . 362
Disabling User Accounts . . . 363
Handling Backups and Auditing . . . 364

Network Address Translation . . . 364
One-to-One NAT . . . 365
Many-to-One NAT . . . 365

Firewall Configuration Example . . . 367

Chapter Summary . . . 372

Key Terms . . . 373

Review Questions . . . 373

Hands-On Projects . . . 375

Case Projects . . . 384

CHAPTER 11
VPN Concepts . . . 385

Understanding VPN Concepts . . . 386
VPN Components . . . 387
Types of VPNs . . . 388
Evaluating Business Needs for VPNs . . . 389
Advantages and Disadvantages of VPNs . . . 391

The Three VPN Core Activities . . . 391
Encapsulation . . . 391
Encryption . . . 398
Authentication . . . 404

Examining VPN Design and Architecture . . . 406
Mesh Topology . . . 406
Star Topology . . . 407
Hybrid Topology . . . 408
VPN Domains . . . 409
Using VPNs with Firewalls . . . 411
Adjusting Packet-Filtering Rules for VPNs . . . 414
Ensuring Client Security . . . 416

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Auditing VPNs and Setting VPN Policies 418
Using VPN Quarantine 418
Logging VPN Activity 419
Auditing Compliance with VPN Policies 419
Guidelines for VPN Policies 420

Chapter Summary 421

Key Terms 422

Review Questions 424

Hands-On Projects 426

Case Projects 435

CHAPTER 12
Internet and World Wide Web Security 437

Examining the Structure of the Internet 438
Understanding the Structure of the Internet 438
Tier System 439
Understanding Weak Points in the Internet's Structure 440

Web Site Attack Techniques 443
Attack Techniques Against Web Servers 443
Buffer Overflow Attacks 443
SQL Injection Attacks 445
Attack Techniques Against Web Users 449

Hardening Web and Internet Resources 453
Hardening DNS Servers 453
DNSSEC 455
Hardening Windows Web Servers 456
Configuring Security Settings in Apache Web Server 458

Chapter Summary 460

Key Terms 461

Review Questions 462

Hands-On Projects 464

Case Projects 474

CHAPTER 13
Security Policy Design and Implementation 475

Understanding the Security Policy Life Cycle 477
Needs Assessment 477
System Design 478
System Implementation 478
Performance Monitoring 478

Examining the Concepts of Risk Analysis 479
Risk Analysis Factors 480
Risk Analysis Methods 484
The Risk Analysis Process 486
Analyzing Economic Impacts 487
Techniques for Minimizing Risk 489

Examining the Concepts of Security Policies 493
General Best Practices for a Security Policy 494
Developing Security Policies from Risk Assessment 495

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Teaching Employees About Acceptable Use 496
Outlining Penalties for Violations 496
Criminal Computer Offenses 496
Enabling Management to Set Priorities 497
Dealing with the Approval Process 498
Feeding Security Information to the Security Policy Team 498
Helping Network Administrators Do Their Jobs 498
Using Security Policies to Conduct Risk Analysis 499

Developing a Security Policy 499
Steps to Creating a Security Policy 500
Identifying Security Policy Categories 501

Defining Incident Handling Procedures 507
Assembling a Response Team 507
Specifying Escalation Procedures 508
Responding to Security Incidents 509
Including Worst-Case Scenarios 509
Updating the Security Policy 510
Conducting Routine Security Reviews 510

Chapter Summary 511

Key Terms 512

Review Questions 514

Hands-On Projects 516

Case Projects 521

CHAPTER 14
Ongoing Security Management 525

Strengthening Control: Security Event Management 526
Monitoring Events 526
Managing Data from Multiple Sensors 528
Evaluating IDPS Signatures 530
Managing Change 531

Strengthening Analysis: Security Auditing 532
Operational Auditing 533
Independent Auditing 534

Strengthening Detection: Managing an IDPS 534
Maintaining Your Current System 534
Changing or Adding Software 535
Changing or Adding Hardware 535

Strengthening Defense: Improving Defense-in-Depth 536
Active Defense-in-Depth 536
Adding Security Layers 537

Strengthening Performance: Keeping Pace with Network Needs 538
Managing Memory 538
Managing Bandwidth 538
Managing Storage 539

Staying Informed About Security Trends 539
Web Sites 539
Mailing Lists and Newsgroups 539
Certifications 540

Chapter Summary 540

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Key Terms 541
Review Questions 542
Hands-On Projects 544
Case Projects 547

APPENDIX A
Security Resources **549**

GLOSSARY **553**

INDEX **567**

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Introduction

This book is intended to provide students and professionals with a solid foundation in the fundamentals of advanced network security. The previous edition of this book placed significant emphasis on intrusion detection, but this edition aims to provide a more balanced approach to the topic of network defense and countermeasures. As the range of threats to data systems becomes broader, depending on a limited number of security strategies becomes riskier. Information security professionals need to have a broad range of knowledge and skills. As a result, the third edition includes topics such as routing security and cryptography, which play an important role in network defense, as well as newer concepts such as IPv6 and unified threat management, which have begun to play a larger role and are expected to become more important in the future.

Intended Audience

Guide to Network Defense and Countermeasures, Third Edition is intended for students and professionals who need hands-on experience with installing routers, firewalls, proxy servers, and intrusion detection and prevention systems (IDPSs) as well as a strong conceptual understanding of routing, packet signature analysis, firewalls, VPNs, intrusion detection and prevention, wireless network security, cryptography, and security policy management. Readers should be familiar with basic networking concepts such as TCP/IP, gateways, routers, and Ethernet standards.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

New to the Third Edition

This edition varies from the second edition in several ways:

- It includes a more balanced approach to network defense and includes new chapter topics.
- It includes new chapters on TCP/IP (Chapter 2), routing fundamentals (Chapter 4), cryptography (Chapter 5), wireless networking and security (Chapters 6 and 7), and Internet security (Chapter 12).
- Hands-on activities have been removed from the body of the chapters to facilitate continuity.
- Hands-on projects have been updated.

Chapter Descriptions

This book has 14 chapters and one appendix as follows:

Chapter 1, "Network Security Fundamentals," provides a review of fundamental security concepts, such as threats to network security, security controls to mitigate the risk of those threats, and the goals of network security.

Chapter 2, "TCP/IP," explains the fundamentals of the TCP/IP network protocol stack, including TCP/IP subprotocols, IP addressing, subnetting, supernetting, variable length subnet masking, and classless interdomain routing. This information provides a foundation for later discussion of packet analysis, such as examination of IP, ICMP, TCP, and UDP headers. The function and structure of IPv6 is addressed in detail.

Chapter 3, "Network Traffic Signatures," introduces students to packet analysis through identification of signatures associated with normal and abnormal traffic. The chapter discusses normal and abnormal findings in TCP, IP, and ICMP packet headers.

Chapter 4, "Routing Fundamentals," discusses the basics of address resolution and router functions, including routing protocols. Both IPv4 and IPv6 routing concepts are discussed. The chapter also covers routing security concepts, including access control lists, authentication, and encrypted router connections.

Chapter 5, "Cryptography," explains cryptographic concepts such as primitives, pseudorandom number generation, hashing, encryption algorithms, digital signatures, Public-key Infrastructure, cryptographic standards, Web security, IPsec, and attacks against cryptography.

Chapter 6, "Wireless Network Fundamentals," discusses concepts of radio frequency transmission, infrared transmission, and signal behavior. The chapter addresses analog and digital modulation along with wireless LANs and wireless standards.

Chapter 7, "Understanding Wireless Security," addresses wireless security concepts and common attacks against wireless networks. The chapter discusses security solutions that are available both for wireless networks and handheld wireless devices. IEEE 802.11 media access control is explained as well.

Chapter 8, "Intrusion Detection and Prevention Systems," identifies the role of IDPSs in network defense; typical detection and prevention methods, including anomaly and signature detection; network and host-based systems; the development of signature rules; and management procedures.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Chapter 9, "Firewalls," provides students with a strong foundation in software and hardware firewalls, with an emphasis on packet filtering and the creation of rule sets.

Chapter 10, "Firewall Design and Management," builds on the previous chapter to address firewall configuration design and proxy server installation and management. Students learn about bastion hosts, honeypots, and Network Address Translation. The chapter also discusses unified threat management concepts and practice.

Chapter 11, "VPN Concepts," presents basic VPN concepts, including encapsulation, encryption, and authentication. VPN configuration and deployment are discussed as well as VPN packet-filtering rules and VPN policies and procedures.

Chapter 12, "Internet and World Wide Web Security," addresses Internet vulnerabilities and the common attacks against these vulnerabilities, including Web server, buffer overflow, SQL injection, ActiveX, and Java Applet attacks. The chapter also discusses security controls, including DNSSEC.

Chapter 13, "Security Policy Design and Implementation," describes the system development life cycle, risk analysis, determination of security controls, security policy concepts, and incident handling procedures.

Chapter 14, "Ongoing Security Management," discusses ways to improve network security through the management of security events. The chapter also addresses auditing and analyzing security procedures and controls as a means of keeping an organization's security posture up to date.

Features of the Book

- **Chapter Objectives**—Each chapter begins with a list of the concepts to be mastered. This list provides a quick reference to the chapter's contents and can be a useful study aid.
- **Chapter Summaries**—Following each chapter discussion is a summary of the concepts introduced in the chapter. These summaries provide students with a quick way to check their understanding of the chapter's main topics.
- **Key Terms**—All terms introduced in boldface text in a chapter are listed and defined after the chapter summary.
- **Review Questions**—The end-of-chapter assessments include a set of questions that allow students to demonstrate their mastery of the chapter's important concepts.
- **Hands-On Projects**—These challenging projects are an important element that gives students an opportunity to practice and research key concepts and skills, and to reinforce the chapter concepts through practical application.
- **Case Projects**—Each chapter contains one or more case projects that provide students with challenging situations for research and analysis.

Text and Graphic Conventions

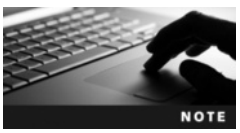

The Note icon draws your attention to additional helpful material related to the subject being discussed.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Tips based on the author's experience provide extra information about how to approach a problem or what to do in real-world situations.

Each hands-on project in this book is preceded by the Hands-On icon and a description of the project.

This icon marks case projects, which are scenario-based assignments. In these extensive case examples, you are asked to implement independently what you have learned.

Instructor Resources

The following supplemental materials are available when this book is used in a classroom setting. All the supplements available with this book are provided to the instructor on a single CD-ROM (ISBN 978-1-1337-2795-8) and online at *www.cengage.com*.

Electronic Instructor's Manual. The Instructor's Manual that accompanies this textbook includes additional instructional material to assist in class preparation, including suggestions for classroom activities, discussion topics, and additional activities.

Solutions. The answers to end-of-chapter material are provided. Solutions are provided for all review questions and for hands-on projects and case projects where applicable.

PowerPoint presentations. This textbook comes with Microsoft PowerPoint slides for each chapter. They are included as a teaching aid for classroom presentation, to make available to students on the network for chapter review, or to be printed for classroom distribution. Instructors, please feel free to add your own slides for additional topics you introduce to the class.

ExamView®. This textbook is accompanied by ExamView, a powerful testing software package that allows instructors to create and administer printed, computer (LAN-based), and Internet exams. ExamView includes hundreds of questions that correspond to the topics covered in this book, enabling students to generate detailed study guides that include page references for further review. The computer-based and Internet testing components allow students to take exams at their computers and save the instructor time by grading each exam automatically.

Figure files. All figures and tables in the textbook are reproduced on the Instructor Resources CD. Like the PowerPoint presentations, they are included as a teaching aid for classroom presentation, to make available to students for review, or to be printed for classroom distribution.

Classroom Setup Guidelines

Most hands-on projects in this book are intended to be performed by pairs of students using two computers: one with Windows Server 2008 R2 Enterprise Edition installed as a domain controller, and one installed with Windows 7 Professional Edition as a member of the domain. Both systems

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

should have the latest updates installed; installation procedures for both operating systems are included later in this section.

Multiple pairs of students can work through the activities in a classroom network environment, or two computers can be connected using a hub or switch. Both students in each pair should work together on each element of the hands-on projects because all tasks are unique and students might not be able to repeat projects with the roles reversed. Alternatively, a single student can work with both systems to perform the hands-on projects.

Several hands-on projects require an additional computer that is set up to run Ubuntu Linux. Specific directions for these setup procedures are provided later in the book as needed.

This section also lists the hardware items and software programs required to complete the hands-on projects in the book. For most of the projects, you need three computers, each with the following features:

Hardware Requirements

- Intel or AMD 64-bit, 1.6 GHz, dual-core processor (4-core, 2.0 GHz or greater recommended)
- At least 2 GB of RAM (4 GB recommended)
- 80-GB hard disk
- Internet access
- DVD-ROM drive
- Super VGA (800 × 600) or higher-resolution monitor
- Video card with 128 MB of RAM and support for DirectX 9 or higher
- Keyboard and mouse or compatible pointing device
- One free USB port (optional)
- One USB wireless adapter (optional)
- One PCI Ethernet network interface card for each PC
- CD-R drive and burning software to create Ubuntu CDs for students

Software Requirements

You need the following operating systems and applications:

- Windows Server 2008 R2 Enterprise Edition
- Windows 7 Professional Edition
- Ubuntu Linux
- ZoneAlarm Free Firewall
- Nmap
- WinPcap
- Sawmill
- Wireshark
- TShark

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Toggit Router Simulator
- Microsoft Word 2010
- TrueCrypt
- inSSIDer
- Snort
- Samba
- Microsoft Forefront Threat Management Gateway 2010
- Apache
- Project Risk Analysis
- Network Asset Tracker

Network Setup

- Each system should be configured with a static IP address, subnet mask, and default gateway that are appropriate for the classroom network and that provide access to the Internet.
- All Windows Server 2008 and Windows 7 systems should be configured with an administrative account that has the username *administrator* and a password of *Pa$$word*.
- The instructor should assign each team of students a domain name of team*x*.net, where *x* is a unique number starting at 1. For example, the domain names should be team1.net, team2.net, and so on.
- Each system should be assigned a hostname based on the system and the team number. For example, the hostnames should be Team1Client.team1.net, Team1Srv.team1.net, Team2Client.team2.net, Team2Srv.team2.net, and so on.
- Although a central instructor server is not required, it may be wise to download required software programs to such a server so that download times are decreased for students and correct versions of the software are available for future classes.

Installing Windows Server 2008 R2

1. Turn on the computer.
2. Insert the Windows Server 2008 R2 Enterprise Edition DVD into the DVD-CD drive.
3. Boot to the DVD.

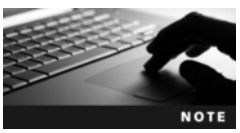

If your system does not boot to the DVD, you might need to alter the device boot order in the BIOS setup utility.

4. In the Install Windows window, verify that the correct language, time, and keyboard type are selected, and click **Next**. Click **Install now**.
5. The next window prompts you to enter your product key for activation. Enter the key number and click **Next**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

6. The next window prompts you to select the operating system you want to install. Click **Windows Server 2008 R2 Enterprise (Full Installation)**, and click **Next**.
7. In the Microsoft Software License Terms window, check the **I accept the license terms** box, and click **Next**. Click **Custom**.
8. The next window asks where you want to install Windows. Click **Drive options (advanced)**, click **New**, enter **30000** in the Size text box, and click **Apply**. Click **OK**. Click **Next**.
9. The system will reboot automatically several times. You are then prompted to change the user password. Click **OK**, enter **Pa$$word** in both text boxes, and press **Enter**. The password is for a user named *administrator* who has full access to the system. Click **OK** in the next window to confirm that your password has been changed.
10. The Initial Configuration Tasks window appears. In the Provide Computer Information section, click **Provide computer name and domain**. In the System Properties window, click the **Change** button. In the Computer Name/Domain Changes window, type **Server*x*** in the Computer name text box, where *x* is the team number assigned by your instructor. Click **OK**. At the prompt that discusses restarting, click **OK**. Click **Close** in the System Properties window, and click **Restart Later** in the Microsoft Windows window.
11. In the Initial Configuration Tasks window, click **Enable automatic updating and feedback** in the Update This Server section. In the Enable Windows Automatic Updating and Feedback window, click **Manually configure settings**. In the Manually Configure Settings window, click the **Change Setting** button in the Windows automatic updating section. In the Change settings window, click **Download updates but let me choose whether to install them** in the drop-down list under Important updates. Click **OK** and then click **Close** in the Manually Configure Settings window.
12. At the bottom of the Initial Configuration Tasks window, check the **Do not show this window at logon** box, and click **Close**.
13. Server Manager opens automatically. In the Server Summary/Computer Information section, check the **Do not show me this console at logon** box, and then close the Server Manager window.
14. Right-click the desktop, click **Screen Resolution**, and choose an appropriate resolution setting for yourself. Close the Screen Resolution window.
15. Click **Start**, and then click **Control Panel**. Type **desktop icons** in the Search Control Panel box. Click **Show or hide common icons on the desktop** under Display. Check the **Computer** and **Network** boxes, and then click **OK** to close the Desktop Icon Settings window. Close the Control Panel.
16. Click the **Start** button on the taskbar, click **Computer**, and double-click **Local Disk (C:)**. From the Organize menu, click **Folder and search Options**. On the View tab in the Folder views section, click **Apply to Folders**, click **Yes**, and click **OK**.
17. Click **Folder Options**, and click the **View** tab. Under Hidden files and folders, click the **Show hidden files, folders and drives** option button, remove the checks from the **Hide extensions for known file types** and **Hide protected operating system files (Recommended)** boxes, read the warning, and click **Yes**. (In a production environment, you should not show hidden files and folders or show protected operating system files on client workstations.) Click **OK** in the Folder Options window.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

18. Click **Start**, and click **Network**. If an information bar appears and informs you that the network discovery and file sharing features are turned off, click the **information bar**, click **Turn on network discovery and file sharing**, and click **Yes, turn on network discovery and file sharing for all public networks**.

The setting described in Step 19 is appropriate in a lab setting, but it should be used with caution and only for a specific business need on a production network. Whenever this information bar appears in a hands-on project in this book, turn on network discovery and file sharing.

19. Click the **Network and Sharing Center** button. In the left pane, click **Change adapter settings**. Right-click **Local Area Connection**, click **Properties**, select **Internet Protocol Version 4 (TCP/IPv4)**, and click the **Properties** button. Click the **Use the following IP address** option button, and then enter the IP address, subnet mask, and default gateway as directed by your instructor. Click **OK** and then click **Close**.
20. Close all windows. Click **Start**, click the right arrow on the far right of the Start menu's bottom line, and click **Shut down**. In the Shut Down Windows window, type **Post-installation reboot** and click **OK**.
21. Upgrade the server to a domain controller using the naming conventions and procedure assigned by the instructor.

Installing Windows 7

1. Turn on the computer.
2. Insert the Windows 7 Professional Edition DVD into the DVD-CD drive.
3. Boot to the DVD.
4. In the Install Windows window, verify that the correct language, time, currency, and keyboard type are selected, and click **Next**. Click **Install now**.

If your system does not boot to the DVD, you might need to alter the device boot order in the BIOS setup utility.

5. In the license terms window, check the **I accept the license terms** box, and click **Next**.
6. When asked which type of installation you want, click **Custom (advanced)** and click **Next**.
7. When asked where you want to install Windows, accept the default location and click **Next**.
8. The system will reboot automatically several times. In the next window, you are prompted to choose a username for your account and to name your computer to distinguish it on the network. Type your first name as the username and type **Win7*x*** as the computer name, where *x* is the team number assigned by your instructor. Click **Next**.
9. In the next window, you set a password for your account. Type **Pa$$word** as the password and enter it again in the second text box. In the Type a password hint text box, type **Pa$$word** again. Note that in a production environment, you would not type the password itself as a hint. Click **Next**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

10. In the next window, enter the product key provided by your instructor. Click **Next**.
11. In the next window, click **Use recommended settings**.
12. In the time and date settings window, verify that the settings are correct and click **Next**.
13. In the window that prompts you to select the computer's current location, click **Work network**.
14. The system opens to the desktop. Right-click the desktop, click **Personalize**, click **Display**, click **Adjust resolution**, and then select a resolution that is appropriate for you. Click **OK** in the Screen Resolution window.
15. After approving the resolution, click **Personalization** in the left pane of the Display window. In the left pane of the Personalization window, click **Change desktop icons**, and check the **Computer** and **Network** boxes. Click **OK** and close the Personalization window.
16. Click the **Start** button on the taskbar, and then click **Control Panel**. Select **Small icons** from the View menu. In the left pane, click **Classic View**. Close the Control Panel.
17. Click **Start**, and then click **Computer**. From the Organize menu, click **Folder and search options**. Click the **View** tab. Under Hidden files and folders, click the **Show hidden files, folders and drives** option button, remove the checks from the **Hide extensions for known file types** and **Hide protected operating system files (Recommended)** boxes, read the warning, and click **Yes**. (In a production environment, you should not show hidden files and folders or show protected operating system files on client workstations.) Click **OK**. Close the Computer window.
18. Right-click the **Start** button, and click **Properties**. Click the **Customize** button. Scroll down, check the **Network** box, and click **OK**. Click **OK** in the next window. Click the **Start** button and click **Network**. If an information bar appears and informs you that the network discovery and file sharing features are turned off, click the **information bar** and then click **Turn on network discovery and file sharing**.

The setting described in Step 18 is appropriate in a lab setting, but it should be used with caution and only for a specific business need on a production network. Whenever this information bar appears in a hands-on project in this book, turn on network discovery and file sharing.

19. Click **Network and Sharing Center** in the menu bar, and click **Change adapter settings**. Right-click **Local Area Connection**, click **Properties**, select **Internet Protocol Version 4 (TCP/IPv4)**, and click the **Properties** button. Click the **Use the following IP address** option button, and then enter the IP address, subnet mask, and default gateway as directed by your instructor. Click **OK**, and then click **Close**.
20. Click the **Start** button, click **All Programs**, and click **Windows Update**. Follow the directions to install all recommended updates.
21. Join the domain created by your partner's server. Use the naming conventions and directions provided by your instructor.
22. Close all windows. Click **Start** and click **Shut down**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Acknowledgments

I would like to thank authors Randy and Dawn Weaver, whose previous editions and vision for this third edition made my job a pleasure. Thanks, too, to the editorial and production staff of Cengage Learning, including Natalie Pashoukos and Andrea Majot, and Suwathiga Velayutham of Integra. In particular, thanks to editor Dan Seiter, whose talent allows me to pretend I'm a writer. I am indebted to reviewer Guy Garrett of Gulf Coast Community College, whose careful analysis and detailed suggestions improved this book significantly.

Thanks to my wife Lisa for letting me spend months of weekends being a desk jockey.

This book is dedicated to my students at Heald College, San Francisco.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 1

Network Security Fundamentals

After reading this chapter and completing the exercises, you will be able to:

- Describe the threats to network security
- Explain the goals of network security
- Describe a layered approach to network defense
- Explain how network security defenses affect your organization

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

This chapter reviews the fundamental network security concepts you need to know. First, you learn about different kinds of intruders and threats to network security, such as threats within your organization, malicious code, and natural disasters. Attackers have many motivations for hacking into networks, and your job is to figure out what they are doing and prevent them from carrying out their plans. The Internet is widely used in most network environments, so you also review some concerns about Internet access.

Next, you learn about the goals of network security and the challenges of ensuring confidentiality, integrity, and availability for network resources. You then delve into the basics of network defense technologies. You discover how layering technologies can ensure better protection than any single technology used alone. The method of layering defensive technologies is called defense in depth (DiD), and includes physical, logical, and virtual security concepts. Auditing is the mainstay of monitoring and troubleshooting a network, so you also review log file basics. Finally, you see how security efforts affect an organization and learn that information security is not the sole domain of information technology (IT).

Examining Network Security Fundamentals

A variety of attackers might attempt network intrusions, causing loss of data, loss of privacy, and other consequences. You learn about these attackers in the following sections. These types of threats are becoming a concern for a growing number of corporate managers. More businesses are actively addressing information security, but many others have not taken steps to secure their systems from attack.

Threats to Network Security

When planning network security measures, knowing the types of attackers who might try to break into your network is important. This knowledge can help you anticipate threats and set up detection systems, firewalls, and other countermeasures to block attacks as effectively as possible. Similarly, understanding the motivation of attackers helps you prepare security controls:

- *Status*—Some attackers attempt to take over computer systems just for the thrill of it. They like to count the number of systems they have accessed as notches on their belt.
- *Revenge*—Disgruntled current or former employees might want to retaliate against an organization for policies or actions they consider wrong. They can sometimes gain entry through an undocumented account (back door) in the system.
- *Financial gain*—Other attackers have financial profit as their goal. Attackers who break into a network can gain access to financial accounts. They can steal individual or corporate credit card numbers and make unauthorized purchases. Just as often, attackers defraud people out of money with scams carried out via e-mail or other means.
- *Industrial espionage*—Proprietary information is often valuable enough that it can be sold to competing companies or other parties.

Hackers A **hacker** is anyone who attempts to gain access to unauthorized resources on a network, usually by finding a way to circumvent passwords, firewalls, or other protective measures. Hackers seek to break into computers for different reasons:

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- "Old school" hackers consider themselves seekers of knowledge; they operate on the theory that knowledge is power, regardless of how they come by that knowledge. They are not out to destroy or harm; they want to discover how things work and open any sources of knowledge they can find. They believe the Internet was intended to be an open environment, and that anything online can and should be available to anyone.
- Other less "ethical" **crackers** pursue destructive aims, such as the proliferation of viruses and worms, much like vandals.
- Some bored young people who are highly adept with computers try to gain control of as many systems as possible for the thrill of it. They enjoy disrupting systems and keeping them from working, and they tend to boast about their exploits online.
- Criminals and industrial spies might be interested in selling information to the top bidder or using it to influence potential victims. Some companies would certainly be interested in getting the plans for a new product from their competitors.
- The term **script kiddie** is often used to describe relatively unskilled programmers who spread viruses and other malicious scripts to exploit weaknesses in computer systems. Script kiddies lack the ability to create viruses or Trojan programs on their own, but they can usually find these programs online.
- **Packet monkeys** are primarily interested in blocking Web site activities through a **distributed denial of service (DDoS) attack**. In a DDoS attack, many computers are hijacked and used to flood the target with so many false requests that the server cannot process them all, and normal traffic is blocked. Packet monkeys might also want to deface Web sites by leaving messages that their friends can read.
- **Hactivists** are computer attackers with political goals. Frequently they use denial of service attacks to shut down Web sites of organizations with whom they disagree. One of the best-known hactivist groups, named Anonymous, has successfully shut down sites of the U.S. Federal Trade Commission to express its opposition to proposed laws that combat digital piracy. Anonymous has also shut down sites that belong to the State of Alabama in protest of immigration laws. After discovering that the Central Intelligence Agency (CIA) was investigating the group, Anonymous shut down some of the CIA's sites as well.

Disgruntled Employees Disgruntled employees are usually unhappy over perceived injustices and want to exact revenge by stealing information. With the economic downturn, more current or former employees are stealing information for financial reasons. Often they give confidential information to new employers. When an employee is terminated, security measures should be taken immediately to ensure that the employee can no longer access the company network and telecommunications systems.

While most attacks come from outside a company, according to CyberSecurity Watch, insider attacks are more costly to a victimized company and are becoming increasingly more sophisticated. Theft, data loss, and network damage can result from the malicious actions of current or former employees. The following are just a few examples:

- A **logic bomb** is malware designed to start at a specific time in the future or when a specified condition exists. At Fannie Mae, the Federal National Mortgage Association,

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

a former engineer planted a logic bomb that could have shut the company down and cost millions by destroying all 4000 of the company's servers. Fortunately, the attack did not succeed. The former employee was sentenced to three years in jail.

- Ansir Khan, a former bank employee in Sheffield, England, attempted to steal $1.9 million after successfully stealing more than $1.1 million from the bank in April 2005 and May 2006. He extracted customer data and shared it with accomplices. He was sentenced to three years in jail.
- A former employee of United Way in Miami, Luis Robert Altamirano, accessed the United Way computer system a year after he left the organization. He deleted files and disabled the voicemail system. Altamirano pled guilty and was sentenced to 18 months in jail and fined $50,000 for computer fraud.
- Adeniyi Adeyemi, a contract employee of Bank of New York Mellon, stole the personal information of dozens of bank employees, mainly in the IT department. He used the information to open dummy financial accounts and receive funds stolen from the accounts of charities and nonprofit organizations.

Terrorists Until September 11, 2001, most people did not consider a terrorist attack on an information infrastructure (known as cyberterrorism) to be a likely threat. Since then, the threat posed by terrorists has been taken more seriously. A terrorist group might want to attack computer systems for several reasons: to make a political statement or accomplish a political goal, such as the release of a jailed comrade; cause damage to critical systems; or disrupt the target's financial stability. Attacking the World Trade Center certainly accomplished the latter goal, given the nature and location of the structures. Terrorists might also want simply to cause panic.

It might be hard to understand why a terrorist attack on computers would be considered a serious threat until you think about how many critical systems are controlled by computers. Consider the chaos that could result from a successful attack on a computer system that controls a nuclear power plant's reactors. The overall psychological effect could be just as detrimental as the infrastructure damage and even the loss of life.

Government Operations The shady world of international espionage is difficult to document, but it is becoming clear that a number of countries see computer operations as more than simply a spying technique; computer networks are a potential battleground. In 2010, a sophisticated malware program called Stuxnet was discovered. The Stuxnet worm was designed to attack Windows systems used in industrial and military settings. The goal was to infect the control systems of automated industrial processes. Security experts who analyzed Stuxnet concluded that it was probably the work of a government operation because of the complexity of the program and the amount of time and resources required to create and propagate it. Because Stuxnet was unusually prevalent in Iran, many observers believe that the United States and/or Israel were responsible for its creation and that it was intended to target Iran's nuclear industry.

Another focus of attention is the Chinese government, which is thought to be responsible for successful computer-based attacks on U.S. Department of Defense information systems as well as government, industrial, and military systems in Germany, France, and Britain.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Malicious Code In 2004, the MyDoom worm infected millions of computers in only a few days, costing $38.5 billion in cleanup, lost productivity, and other losses. MyDoom was believed to have been the fastest-spreading worm ever created. MyDoom is primarily transmitted via e-mail, with subject lines such as "Error," "Mail Delivery System," or "Mail Transaction Failed." If the user opens the attachment, the worm resends itself to e-mail addresses in the user's address book and local files. The first variant, MyDoom.A, contained a back door on port 3127/tcp and a denial of service attack on the SCO Group Web site that was timed to launch on February 1, 2004. The second variant, MyDoom.B, targeted the Microsoft Web site. It blocked access to Microsoft and some online antivirus sites, thus denying access to antivirus updates and virus-removal tools.

In 2008, a worm known as Conficker was discovered. This program attacked all Windows operating systems from Windows 2000 through Windows 7. An estimated 9 to 15 million computers were infected. In 2009, Microsoft offered a $250,000 reward for the identification of Conficker's authors. Conficker was designed to create **botnets**: networks of tens of thousands of infected computers that belong to unsuspecting victims and can be controlled from a central station. As of this writing, the authors of Conficker have not been identified, but because the program was designed not to infect systems with a Ukrainian keyboard, it is thought that the worm was developed in Eastern Europe.

Information security has improved since MyDoom and Conficker, but new vulnerabilities always lurk right around the corner, and security professionals must stay one step ahead of attackers. The following sections review the types of malware you might encounter.

Viruses, Worms, and Trojan Programs Although most users think of any type of virus, worm, or Trojan program as similar problems, they are completely different types of attacks. A **virus** is executable code that can replicate itself from one place to another surreptitiously and perform actions that range from benign to harmful. Viruses are spread by several methods, including running executable code, sharing disks or memory sticks, opening e-mail attachments, and viewing infected or malicious Web pages. Viruses can attach to other executables or replace them in order to spread or execute. Viruses require user intervention to run.

A **worm** creates files that copy themselves repeatedly and consume disk space. Worms do not require user intervention to be launched; they are self-propagating. Some worms can install **back doors**—a way of gaining unauthorized access to a computer or other resource, such as an unused port or terminal service, that makes it possible for attackers to gain control over the computer. A **port** is an area in random access memory (RAM) that is assigned a number (the port address) and is reserved for a program that runs in the background to listen for requests for the service it offers. Other worms can destroy data on a hard disk. Just like a cold or flu virus, computer viruses and worms can mutate or be altered to defeat antivirus software.

A **Trojan program** is also a harmful computer program, but one that appears to be something useful—a deception like the Trojan horse described in Greek legends. The difference between a virus and a Trojan program lies in how the malicious code is used. Viruses replicate and can potentially cause damage when they run on a user's

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

computer. Trojan programs can also create a back door, which opens the system to additional attacks. The often hidden or obscure nature of a back door makes the attacker's activities difficult to detect.

Viruses, worms, and Trojan programs are a major security threat. They can damage files, enable attackers to control computers, and prevent applications from functioning correctly. When creating a network defense perimeter, you need to consider guarding against all three. Firewalls and intrusion-detection systems do not block malicious code on their own; you need to install anti-malware software or proxy servers that can be configured to filter out malicious code and delete it before it causes harm.

Macro Viruses A **macro** is a type of script that automates repetitive tasks in Microsoft Word or similar applications. When you run a macro, a series of actions are carried out automatically. Macros are a useful way to make some tasks perform more efficiently. Unfortunately, macro viruses perform the same functions as macros, but they tend to be harmful. For example, in March 1999, the Melissa macro virus caused Microsoft to shut down the company's entire e-mail service. Melissa spread rapidly and arrived as an attachment with the subject line "Important message from [*name of someone*]." The body text read, "Here is that document you asked for…don't show anyone else." If the recipient opened the attachment, the macro virus infected the computer and carried out a series of commands. Melissa was a fast-spreading virus, infecting more than 100,000 computers in the first few days. Macro viruses remain a threat today, but the good news is that the user must perform some action for the virus to be activated; therefore, educating users not to open the attachments is essential. Most modern operating systems and office suites do not automatically run macros, so the threat from macro viruses is reduced.

Other Threats to Network Security You cannot prepare for every possible risk to your systems. At best, you can maintain a secure environment for today's threats and have a comprehensive plan for integrating safeguards against tomorrow's threats into your defenses. The next threat might be infection by a new virus, exploitation of a recently discovered vulnerability, or an earthquake that destroys your facility. Many threats, such as natural disasters, cannot be mitigated entirely. Although you might have prepared for natural disasters by maintaining an alternate site complete with all necessary equipment, your primary site's network and equipment could still suffer devastating loss.

Social Engineering: The People Factor One common way that attackers gain access to an organization's resources cannot be prevented with hardware or software. The vulnerability in this case is well-meaning but gullible employees who attackers fool into giving out passwords or other access codes. To protect itself against personnel who do not observe accepted security practices or who willfully abuse them, an organization needs a strong and consistently enforced security policy and a rigorous training program. Security policies are discussed in Chapter 13.

Common Attacks and Defenses

Table 1-1 describes some of the common attacks you need to guard against and the defensive strategies you can use to defeat them. These concepts are discussed in more depth throughout the remainder of the book.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Attack	Description	Defense
Denial of service (DoS) attack	The traffic into and out of a network is blocked when servers are flooded with malformed packets (bits of digital information) that contain false IP addresses, other harmful data, or other fake communications.	Keep your server OS up to date; log instances of frequent connection attempts against one service.
SYN flood	A network is overloaded with packets that have the SYN flag set. Servers are overloaded with requests for connections and are unable to respond to legitimate requests (a denial of service attack).	Keep your firewall and OS up to date so that these attacks are blocked by means of software patches and updates, and review your log files of access attempts to see whether intrusion attempts have been made.
Virus	Network computers are infected by viruses.	Install antivirus software and keep virus definitions up to date. Keep applications and operating systems patched.
Trojan program	A user installs a malicious Trojan program that can create a "back door" an attacker can exploit.	Install antivirus software and keep virus definitions up to date. Keep applications and operating systems patched.
Social engineering	An employee is misled into giving out passwords or other sensitive information.	Educate employees about your security policy, which is a set of goals and procedures for making an organization's network secure.
Malicious port scanning	An attacker looks for open ports to infiltrate a network.	Install and configure a firewall, which is hardware and/or software designed to filter out unwanted network traffic and protect authorized traffic.
Internet Control Message Protocol (ICMP) message abuse	A network is flooded with a stream of ICMP echo requests to a target computer.	Set up packet filtering.
Man-in-the-middle attack	An attacker operates between two computers in a network and impersonates one computer to intercept communications.	Use VPN encryption.
Finding vulnerable hosts on the internal network to attack	An attacker who gains access to one computer on a network can get IP addresses, host names, and passwords, which are then used to find other hosts to attack.	Use proxy servers.
New files being placed on the system	A virus or other program causes new files to proliferate on infected computers, using up system resources.	Install system-auditing software.
Remote Procedure Calls (RPC) attacks	The operating systems crash because they are unable to handle arbitrary data sent to an RPC port.	Set up an IDPS (intrusion detection and prevention system).

© Cengage Learning 2014

Table 1-1 Common attacks and defenses (*continues*)

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Attack	Description	Defense
Application vulnerability exploits	Unpatched or vulnerable client-side applications that can be invoked and then misused by browsers are targeted, often by trusted Web sites converted into malicious servers.	Keep applications patched. Maintain software inventories so vulnerable software is accounted for and defended. Ensure secure configurations of all software and use perimeter defenses to help identify and prevent attacks.
Web application attacks	Brute force password guessing is used to gain a valid username/password pair. Popular targets of this attack are Microsoft SQL, SSH servers, and FTP. Cross-site scripting, SQL injection, and PHP File Include attacks are the most popular methods for compromising Web sites.	Perimeter defenses should be used to ensure that layered defenses identify and prevent attacks aimed at Web servers. Log files can help determine if your Web server has been compromised. Ensure that all applications and operating systems are patched regularly.

© Cengage Learning 2014

Table 1-1 Common attacks and defenses (*continued*)

Internet Security Concerns As you probably know from your study of basic networking concepts and TCP/IP, a port number combined with a computer's IP address constitutes a network connection called a **socket**. Attackers commonly use software to try to identify sockets that respond to connection requests. The sockets that respond can be targeted to see whether they have been left open or have security vulnerabilities that can be exploited. Hypertext Transport Protocol (HTTP) Web services use port 80. HTTP is among the most commonly exploited services. In Hands-On Project 1-3, you use a port scanning tool to test the ports of a target computer.

E-Mail and Communications For a home user who regularly surfs the Web, uses e-mail, and engages in instant messaging, a firewall's primary job is to keep viruses from infecting the system and to prevent Trojan programs from being installed and creating back door openings. Personal firewall programs, such as Comodo Internet Security, come with an antivirus program that alerts users to an e-mail attachment or a file containing a known virus.

Scripting A widespread network intrusion that is increasing in frequency and severity is the use of **scripts**—executable code attached to e-mail messages or downloaded files that infiltrates a system. It can be difficult for a firewall or intrusion-detection system (IDS) to block all such files; specialty firewalls and other programs should be integrated with existing security systems to keep scripts from infecting a network.

A specialty e-mail firewall can monitor and control certain types of content that pass into and out of a network. These firewalls can be configured to filter out pornographic content, junk e-mail, and malicious code. M86 MailMarshal Secure Email Gateway by M86 Security, for instance, scans the content of each e-mail message before it reaches the recipient.

E-mail filtering programs, however, introduce privacy issues that need to be balanced against an organization's need for protection—a trade-off that applies to almost all aspects of network security, not just e-mail messages. Another problem you might encounter with

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

the glut of available security software is redundant program functions that are incompatible. For example, Windows Firewall is included in Windows 7; if you run it and another personal firewall program at the same time, you may have problems with your connection and desirable programs or communications being blocked.

Always-On Connectivity The proliferation of affordable high-speed connections, such as cable and DSL, brings up special security concerns for network administrators. Computers that use always-on connections are easier to locate and attack because their IP addresses remain the same as long as they are connected to the Internet—which might be days at a time if computers are left on overnight or over a weekend. Some users pay extra for static IP addresses that never change and that enable them to run easily found Web servers or other services. Static IP addresses, however, make it easier for attackers to locate a computer and scan it for open ports.

Another problem could occur when remote users want to connect to an organization's internal network. Remote users include employees who are on the road, contractors who work at home, and business partners. As the Internet grew in popularity, more home computers started using modems. These connections were usually made through temporary dial-up connections that used protocols such as Point-to-Point Protocol (PPP). Today, however, it is increasingly likely that remote users connect to a network through an always-on DSL or cable modem connection, which means they might be connected to a network for hours at a time.

Always-on connections effectively extend the boundaries of your corporate network, and you should secure them as you would any part of your network perimeter. At the very least, your network security policy should specify that remote users have their computers equipped with firewall and antivirus protection software. While a written policy may be helpful, more and more organizations are using technology to enforce such remote access policies and allow administrators to block internal network access to remote systems that have not met security requirements. After all, if attackers can break into a remote user's computer while the user is connected to your network through a virtual private network (VPN) or other connection, your network becomes vulnerable as well.

Goals of Network Security

Providing Secure Connectivity In the early days of the Internet, network security primarily emphasized blocking attackers and other unauthorized users from accessing the corporate network. Today the priority is secure connectivity with trusted users and networks. When people go online to conduct business, they often engage in activities that could make them vulnerable:

- Placing orders for merchandise online, revealing both personal and financial information during payment
- Paying bills by transferring funds online
- Accessing account information
- Looking up personnel records
- Creating authentication information, such as usernames and passwords

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The growth of the Internet and e-commerce is not likely to slow down, so methods to secure these transactions must be set up and maintained. Several methods can be combined in a layered security scheme, as you learn in the next section.

Secure Remote Access One of the biggest security challenges for organizations that communicate via the Internet is the need to provide secure remote access for contractors and employees who are traveling. A VPN, with its combination of encryption and authentication, is often provided by the industry standard, IP security (IPsec), and is a simple and cost-effective solution (see Figure 1-1). VPNs are explained in more detail in Chapter 11.

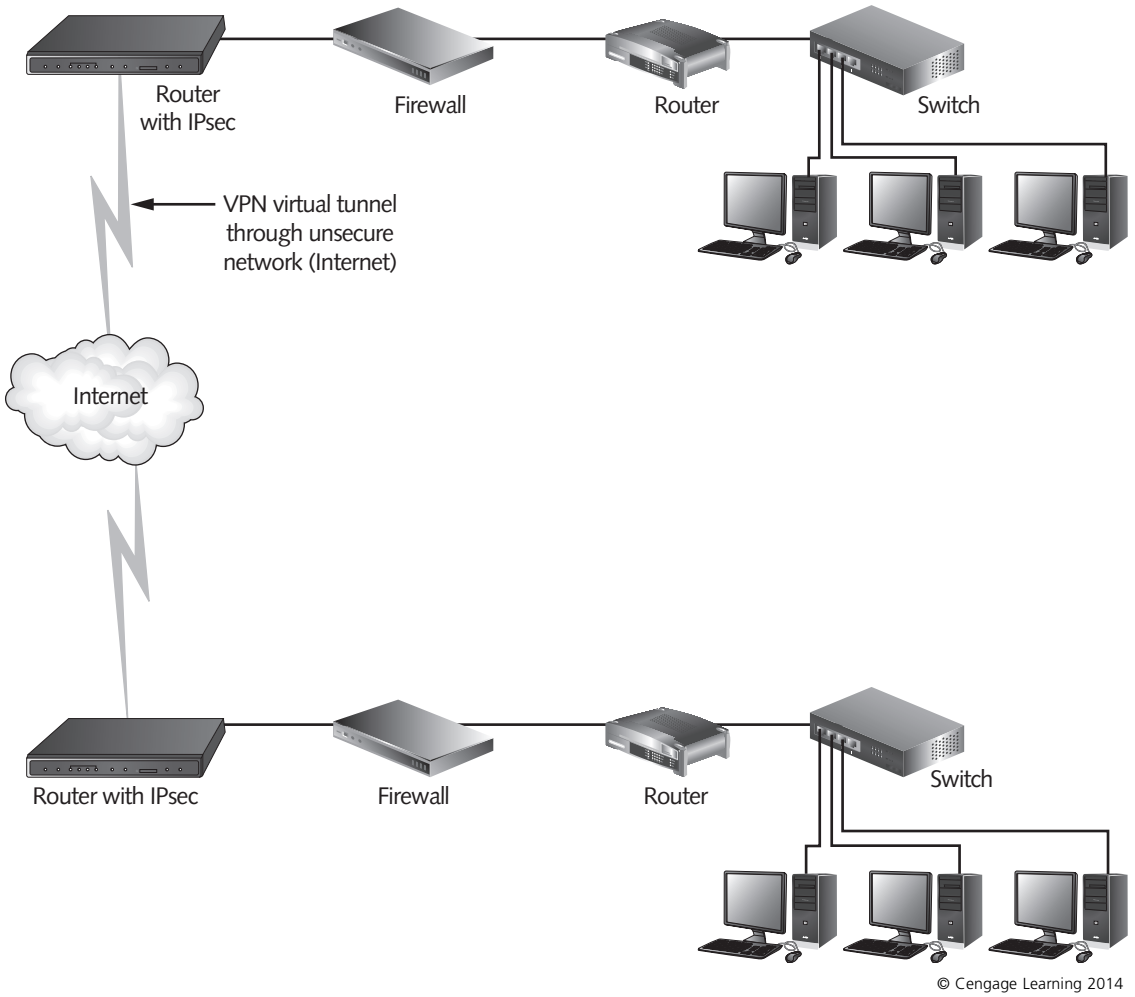

© Cengage Learning 2014

Figure 1-1 Many businesses provide secure remote access using VPNs

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Ensuring Privacy Corporations, hospitals, and other organizations that maintain databases of personal and financial information need to maintain privacy, not only to protect their customers but to maintain the integrity and credibility of their own companies. In addition, U.S. law protects private information and mandates severe penalties for failure to protect it adequately. Examples of these laws include Sarbanes-Oxley (SOX) for publicly traded companies, the Health Insurance Portability and Accountability Act (HIPAA) for medical organizations, the Family Educational Rights and Privacy Act (FERPA) in education, the Payment Card Industry Data Security Standard (PCI DSS) for organizations that accept credit cards, and the Gramm-Leach-Bliley Act in banking. If you work in an industry affected by laws that govern privacy protection, you may have a compliance department that can help you keep current with legal requirements.

One of the most important and effective ways to maintain the privacy of an organization's network is to educate all employees about security dangers and to explain security policies. Employees are likely to detect security breaches *and* to cause security breaches accidentally through their own behavior. They can also be mindful of their coworkers' activities and be alert to suspicious actions that could indicate a security problem.

Providing Nonrepudiation Encryption protects the integrity, confidentiality, and authenticity of digital information. Encryption can also provide **nonrepudiation**, which is the capability to prevent a participant in an electronic transaction from denying that it performed an action. Nonrepudiation simply means ensuring that the sender cannot deny sending a message and the recipient cannot deny receiving it. Nonrepudiation is an important aspect of establishing trusted communication between organizations that do business across a network rather than face to face.

Confidentiality, Integrity, and Availability The three primary goals of information security are data confidentiality, data integrity, and data availability. It is hard to imagine any aspect of information technology that has no responsibility for ensuring one or more of these three fundamental goals. **Confidentiality** is the prevention of intentional or unintentional disclosure of communications between a sender and recipient. **Integrity** ensures the accuracy and consistency of information during all processing (creation, storage, and transmission). **Availability** is the assurance that authorized users can access resources in a reliable and timely manner.

Using a Layered Defense Strategy: Defense in Depth

At this comparatively early stage in the development of information security, no single security component or method can be expected to ensure complete protection for a network—or even an individual host computer. Instead, you need to assemble a group of methods that work in a coordinated fashion to provide protection against a variety of threats. Even then, it is not realistic to think that all security threats can be stopped. Security is more a state of mind than a tangible, absolute state.

The components and approaches described throughout the rest of this book should be arranged to provide layers of network defense. This layering approach to network security is

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

often called **defense in depth (DiD)**. The National Security Agency (NSA) originally designed DiD as a best practices strategy for achieving information assurance.

When beginning with an unprotected system, the first layer of defense added is always the most effective. As more layers are stacked on the first, potential attackers must successfully breach each layer to gain access to the next one. However, adding layers also adds increasing complexity for system administrators. Security enhancements must be balanced against the cost to maintain and monitor defenses. DiD does eventually reach a point where the cost of implementing additional security outweighs the potential benefits.

Another goal of implementing DiD should be to find ways that the security layers can work together, each using data generated by others to enhance the overall effectiveness of the systems.

In general, the layers are as follows. Each layer is discussed in the following sections.

- Physical security
- Authentication and password security
- Operating system security
- Antivirus protection
- Packet filtering
- Firewalls
- Demilitarized zone (DMZ)
- Intrusion detection and prevention system (IDPS)
- Virtual private networks (VPNs)
- Network auditing and log files
- Routing and access control methods

Physical Security

The term **physical security** refers to measures taken to protect a computer or other network device from theft, fire, or environmental disaster. Examples of physical security include installing computer locks that attach a computer device to a piece of furniture in your office, and keeping critical servers in a room protected by a lock and/or burglar alarm. *If the bad guys can touch it, they own it*. This statement means that a computer can easily be compromised if a malicious intruder has physical access to it. Within minutes, an attacker can defeat most common physical locks and steal anything from a password file to the whole server. More insidiously, attackers can plant malware that could give them control of the system without the owner's knowledge.

Authentication and Password Security

After you have physically secured your computers, you can begin to protect them from the inside as well. One simple but reasonably effective strategy is a password security policy, which requires your employees to select good passwords, keep them secure, and change them regularly. Using multiple passwords, including screen-saver passwords and passwords for protecting critical applications, is also a good idea to guard against unauthorized employees gaining control of unattended computers. But, unless password policies are in place to

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

ensure the use of complex passwords and their safekeeping can be enforced through technical means, passwords can become a serious vulnerability. These days, more stringent methods of authentication are becoming common.

Authentication—verifying the identity of a user, service, or computer—uses one of three methods: verifying something the user knows, something the user possesses, or something the user is. In the field of network computing, authentication is performed in one of several ways. **Basic authentication** involves using something the user knows, such as a username/password pair. In **challenge/response authentication**, the authenticating device generates a random code or number (the challenge) and sends it to the user who wants to be authenticated. The user resubmits the number or code and adds a secret PIN or password (the response), or uses a possession such as a smart card to swipe through a card reader.

In large organizations, a centralized server typically handles authentication. The use of **biometrics**—physical information that identifies a person, such as retinal scans, voiceprints, and fingerprints—is growing in popularity because of the security limitations of relying on username and password combinations alone.

Operating System Security

Another way to secure computers and their data from the inside is by installing operating system (OS) patches that have been issued to address security flaws. It is your responsibility to keep up with patches, hot fixes, and service packs and to test and install them when they become available. In addition, stopping any unneeded services and disabling guest user accounts helps make an OS more secure.

Antivirus Protection

Virus scanning refers to the process of examining files or e-mail messages for filenames, file extensions such as .exe (for executable code) or .zip (for zipped files), and other indications that viruses are present. Many viruses have suspicious file extensions, but some seem innocuous. Antivirus software uses several methods to look for malware, including comparisons to the software's current **signature files**, which contain a pattern of known viruses. Signature files are the primary reason for keeping your antivirus software updated; antivirus software vendors frequently create updates and make them available for customers to download. When antivirus software recognizes the presence of viruses, it deletes them from the file system or places them in a storage area called a *quarantine* where they cannot replicate themselves or do harm to other files.

Firewalls and IDSs, by themselves, are not equipped to scan for viruses and eliminate them. However, many enterprise-level firewalls come with integrated antivirus protection. Antivirus software is a must-have for every computer in a network; if your firewall does not provide antivirus software, you need to install it on the computer that hosts the firewall and on all network computers.

Packet Filtering

Packet filters block or allow the transmission of packets of information based on port, IP address, protocol, or other criteria. Packet filtering can be performed by different types of systems. Some are hardware devices, such as routers and firewalls placed at a network

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

gateway. Others are software programs that can be installed on a gateway or a computer. Here are a few examples:

- *Routers*—These devices are probably the most common packet filters. Routers process packets according to an access control list (ACL) the administrator defines.
- *Operating systems*—Some systems, such as Windows and Linux, have built-in utilities for packet filtering on the TCP/IP stack of the server software. Linux has a kernel-level packet filter called Iptables; Windows Server 2008 and Windows 7 have a feature called Windows Filtering Platform.
- *Software firewalls*—Most enterprise-level programs, such as Check Point Firewall Software Blade, perform packet filtering. Personal firewalls such as ZoneAlarm perform basic stateless packet filtering based on simple rules. You use ZoneAlarm in the hands-on projects at the end of this chapter.

Whatever type you use, the packet-filtering device evaluates information in the packet header and compares it to the established rules. If the information corresponds to one of the "allow" rules, the packet is allowed to pass; if the information matches one of the "deny" rules, the packet is dropped.

Firewalls

The foundation for installing and configuring a firewall is your organization's overall security policy. After you have a solid security policy as your guide, you can design security configurations to support your organization's goals. Specifically, you can create a packet-filtering rule base for your firewall that reflects your overall approach to network security. (You learn about security policies in Chapter 13 and about firewalls in more detail in Chapters 9 and 10.) The following section describes two ways that a firewall can control the amount of protection a network receives: permissive versus restrictive policies.

Permissive vs. Restrictive Policies A firewall, following the direction given in a security policy, typically adopts one of the following general approaches to security (see Figure 1-2):

- ***Permissive policy***—Calls for a firewall and associated security components to allow all traffic through the network gateway by default and then to block services on a case-by-case basis.
- ***Restrictive policy***—Calls for a firewall and associated network security components to deny all traffic by default. The first rule denies all traffic on any service and using any port. To allow a specific type of traffic, a new rule must be placed ahead of the "deny all" rule.

A firewall should enforce the overall policy established by the network administrator. Enforcement is handled primarily through setting up packet-filtering rules; a rule base contains a set of these rules. The order of rules in the rule base is important to how the firewall processes traffic.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

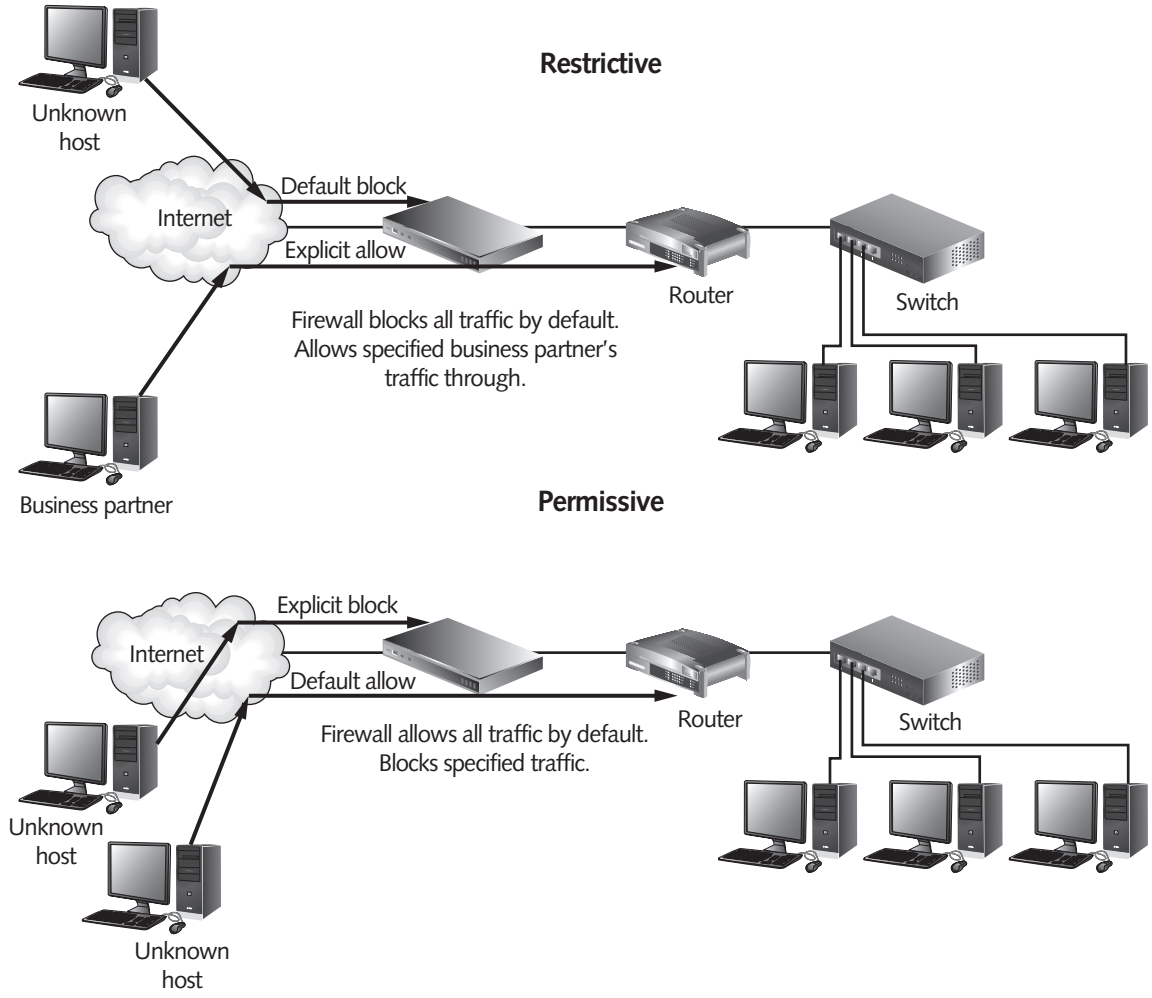

© Cengage Learning 2014

Figure 1-2 Permissive vs. restrictive firewall policies

Demilitarized Zone (DMZ)

A subnet called a **demilitarized zone** (**DMZ**) is a network that sits outside the internal network but is connected to the firewall. A DMZ makes services like HTTP (Web server) and FTP (File Transfer Protocol) publicly available, yet protects the internal LAN. A DMZ might also contain a DNS server that resolves fully qualified domain names to IP addresses. The subnet attached to the firewall and contained in the DMZ is sometimes called a *service network* or *perimeter network*. A common DMZ configuration is shown in Figure 1-3.

Intrusion Detection and Prevention System (IDPS)

Ideally, firewalls and proxy servers block intruders or malicious code from entering a network. However, using an IDPS with these tools offers an additional layer of protection for a network. An **intrusion detection and prevention system** (**IDPS**) works by recognizing the

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

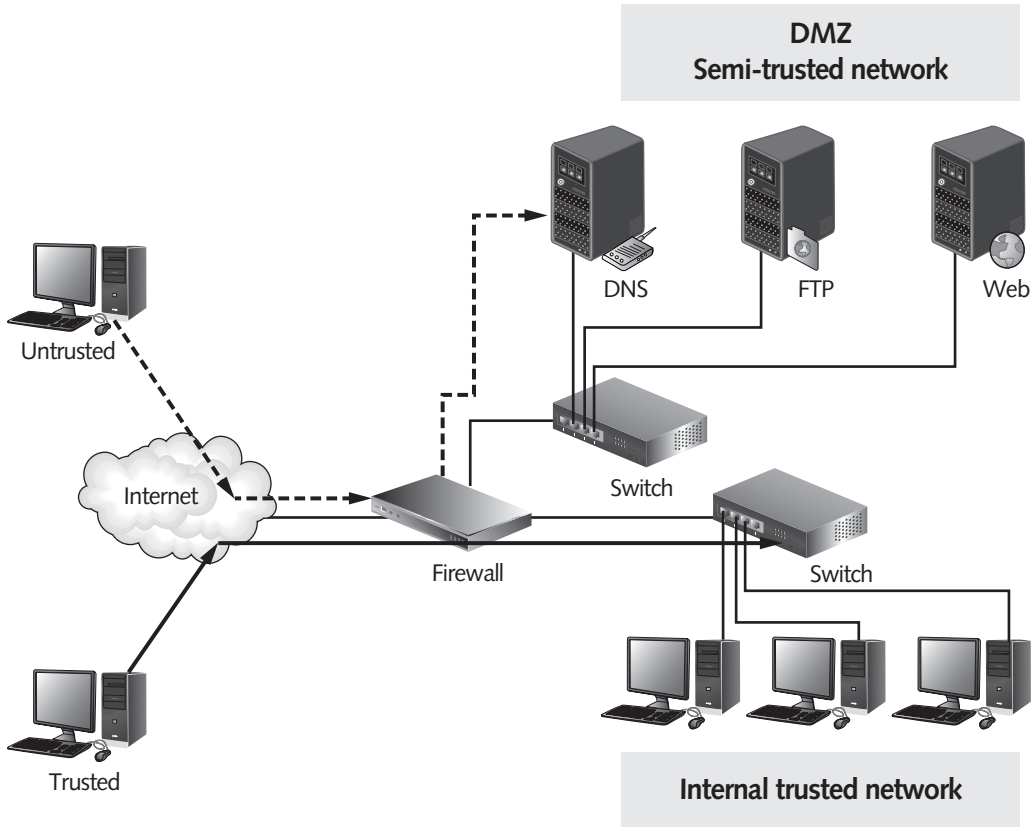

©Cengage Learning 2014

Figure 1-3 Firewall used to create a DMZ and protect the internal network

signs of a possible attack and sending a notification to an administrator that an attack is under way (intrusion detection). Some traffic can trigger a response that attempts to actively combat the threat (intrusion prevention). Note that the term *intrusion prevention* is not precise because there is no known method for preventing all possible intrusions. The signs of possible attacks are commonly called **signatures**—combinations of IP addresses, port numbers, and the frequency of access attempts. You learn the details of IDPS concepts and implementation in Chapter 8.

Virtual Private Networks (VPNs)

Companies that share files or exchange confidential financial information traditionally used expensive leased lines provided by telecommunications companies. Although these lines created a point-to-point connection between company networks and therefore ensured a high level of security, the monthly costs were excessively high for many budget-conscious companies. Today, a more common approach to protecting confidential data in transit is the use of VPNs, which provide a low-cost and secure connection that uses the public Internet.

A **virtual private network** (**VPN**) is a network that uses public telecommunications infrastructure, such as the Internet, to provide secure access to corporate assets for remote users. VPNs

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

use authentication to verify users' identities and encrypt and encapsulate traffic to protect it in transit. VPNs are covered in detail in Chapter 11.

Network Auditing and Log Files

Auditing is the process of recording which computers are accessing a network and what resources are being accessed, and then recording the information in a log file. IT managers often overlook detailed reviews of log files generated by firewalls and IDPS. By reviewing and maintaining log files, you can detect suspicious patterns of activity, such as regular and unsuccessful connection attempts that occur at the same time each day. You can identify those who have attacked your network, or at least gather enough information to begin to identify them. You can set up rules to block attacks and keep your network defense systems up to date by examining attack attempts that have penetrated firewalls and other protective devices. Effective management of log files is an essential activity that goes hand in hand with any perimeter security configuration.

Log File Analysis

- Compiling, reviewing, and analyzing log files are among the most tedious and time-consuming tasks associated with network security. Network administrators read and analyze log files to see who is accessing their networks from the Internet. All connection attempts that are rejected should be recorded in the hope of identifying possible intruders or pinpointing vulnerable points in the system.
- When you first install intrusion detection or firewalls on your network, you will probably be asked to prepare reports that describe how the network is being used and what kinds of filtering activities the device is performing. It is a good idea to sort logs by time of day and per hour. (Sorting log files produces material that is more organized and easier to review than the log files produced by the server, firewall, or other device.)
- Be sure to check logs to learn the peak traffic times on your network, and try to identify the services that consume the largest part of your available bandwidth. If your firewall or IDPS can display log file entries graphically (as shown in Figure 1-4), you should consider showing the graphs to management as needed. Graphs illustrate trends more effectively than lists of raw data.

If your firewall or IDPS cannot display log files graphically, it is well worth the time to locate and install a compatible product that can.

Configuring Log Files Typically, the log files compiled by a firewall or IDPS give you different options. You can view active data (data compiled by the firewall as traffic moves through the gateway in real time) or data that the device has recently recorded. You can also view the information in the following ways:

- *System events*—These events usually track the operations of the firewall or IDPS, making a log entry whenever it starts or shuts down.
- *Security events*—These events are records of any alerts the firewall or IDPS has issued.
- *Traffic*—This is a record of the traffic that passes through the firewall.
- *Packets*—Some programs enable you to view information about packets that pass through them.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

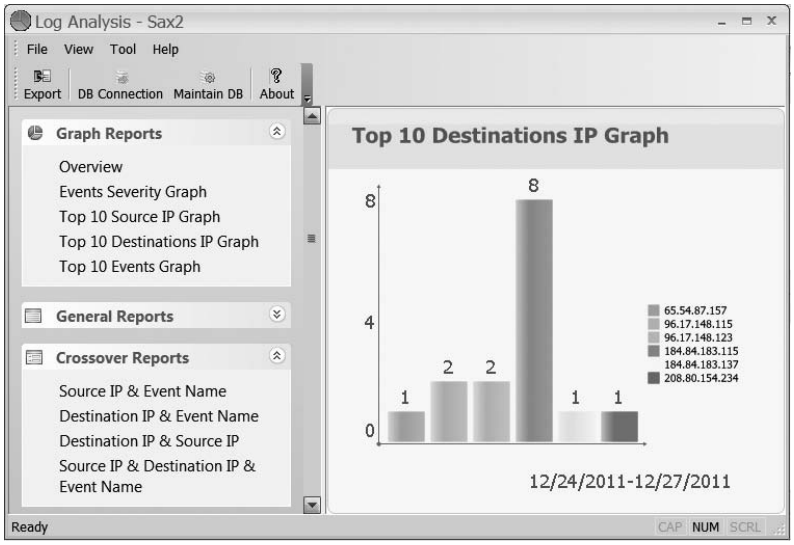

Source: Sax2

Figure 1-4 Graphic display of log file entries

With more elaborate programs, you can customize what you see in log files and search for specific items or events.

Routing and Access Control Methods

Routers at the perimeter of a network are critical to the movement of all network traffic, regardless of whether the traffic is legitimate or harmful. Because routers are positioned on a network's perimeter, they can be equipped with their own firewall software to perform packet filtering and other functions.

To set up a defense, you need to know what kinds of attacks to expect and which of your services and computers might present openings that can be exploited. As a security professional, your goal is ensuring that no unauthorized access occurs. You must identify areas that would allow attackers to gain access to your network. An attacker might attempt to access the following open points of entry:

- *Vulnerable services*—The attacker might be able to exploit known vulnerabilities in an application.
- *E-mail gateways*—The attacker might be able to attach a virus payload to an e-mail message. If a recipient clicks the attachment to open it, the program runs and the virus installs itself on the user's system.
- *Porous borders*—Computers on the network might be listening (that is, waiting for connections) on a port that has no functional use. If an attacker discovers a port that the computer has left open and that is not being used, the open port can give the attacker access to that computer's contents.

Users must have access to the resources necessary to do their jobs, but unauthorized people must not be able to gain access to those resources. Access control is a vital facet of

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

network security and encompasses everything from complex permission configurations on domain controllers to locked doors. You should know about three main methods of access control:

- ***Mandatory access control (MAC)***—This is an uncompromising method for controlling how information can be accessed. With the MAC method, all access capabilities are defined in advance. System administrators establish what information users can and cannot share.
- ***Discretionary access control (DAC)***—With this method, network users are given more flexibility in accessing information. This method allows users to share information with other users; however, the risk of unauthorized disclosure is higher than with the MAC method.
- ***Role-based access control (RBAC)***—This method establishes organizational roles to control access to information. The RBAC method limits access by job function or job responsibility. An employee could have one or more roles that allow access to specific information.

The Impact of Defense

Although the cost of securing systems and their data might seem high, in terms of return on investment (ROI) the cost of a security breach can be much higher. As mentioned, several laws exist to protect privacy, and violation of those laws can carry severe monetary penalties. When added to the direct and indirect costs of a security breach, implementing a sound security scheme can seem inexpensive by comparison.

A key factor in securing systems successfully is the support you gain from upper management. Before security efforts ever start, executives and managers have to be sold on the idea. This support serves several key purposes:

- The project will cost money, and you need to have funding for the project approved beforehand.
- The project will require IT staff time, and managers, supervisors, and employees from all departments must participate to paint a clear picture of priorities and carry out the security plan.
- The process of implementing security systems might require down time for the network, which translates into lost productivity and inconvenience for everyone.
- Most importantly for the long-term success of security efforts, executives and management need to support the project from start to finish. If they do not, development, testing, implementation, and maintenance are nearly impossible to complete. The necessary resources and enforcement will not be available. Besides, if management does not seem to care and does not support the initiative, why would anyone else?

In addition, remember that it is not enough simply to plan and implement security systems. Probably the most challenging facet of information security is keeping up to date on new threats and other developments in the industry. Security systems must be maintained and updated continuously to provide protection against new threats.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Chapter Summary

- Network intruders might be motivated simply by a desire to see what kind of data is available on a network and to gain control of computers. Revenge by disgruntled current or former employees might be the primary motivation, however. Some attackers break into accounts and networks for financial gain. Others want to steal proprietary information for their own use or for resale to other parties.
- Because the Internet plays an important role in moving business-related traffic from one corporate network to another, an understanding of network security concerns pertaining to online communication is essential. E-mail is one of the most important services to secure because of the possibility of malware in e-mail attachments. In addition, always-on connections present security risks that need to be addressed with firewall and VPN solutions.
- Goals for a network security program come from an analysis of risks and an assessment of resources you want to protect. An important goal of any network security effort is maintaining the privacy of customer and employee information. Other goals include preserving data integrity, authenticating approved users of network resources, and enabling remote users to connect to the internal network securely.
- An effective network security strategy involves many layers of defense working together to prevent security threats.
- Auditing is the process of creating a log file and then using it to record information about which computers access a network and which resources are accessed. Review firewall, packet-filtering, and IDPS logs regularly to detect vulnerable points that should be closed.
- Routers at the network perimeter are critical to the movement of traffic into and out of the network. Access control ensures that users can access resources they need but that unauthorized people cannot access network resources to exploit them.
- Defense affects the entire organization. Before an organization begins a security project, management must agree on and support the project. IT staff need input from managers, supervisors, and employees in all departments to create an effective policy and carry out security measures.

Key Terms

auditing The process of recording which computers are accessing a network and what resources are being accessed, and then recording the information in a log file.

authentication The process of verifying the identity of a user, computer, or service.

availability The assurance that authorized users can access resources in a reliable and timely manner.

back doors Ways of gaining unauthorized access to a computer or other resource, such as an unused port or terminal service.

basic authentication An authentication method that uses a username/password pair to verify the identity of the user requesting access.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

biometrics A method of authenticating a user using physical information, such as retinal scans, fingerprints, or voiceprints.

botnets Networks of computers owned by unsuspecting victims of exploitation and controlled from a central system.

challenge/response authentication An authentication method in which one party presents a question, called the challenge, and the other party must provide the correct response, usually a password, to be granted access.

confidentiality Preventing intentional or unintentional disclosure of data during its creation, transmission, and storage.

crackers Hackers who break into systems with the intent of doing harm or destroying data.

defense in depth (DiD) A strategy for achieving information security that uses multiple layers of defense.

demilitarized zone (DMZ) A semitrusted subnet that lies outside the trusted internal network but is connected to the firewall to make services publicly available while still protecting the internal LAN.

discretionary access control (DAC) An access control method that allows users to share information with other users; however, the risk of unauthorized disclosure is higher than with the MAC method.

distributed denial of service (DDoS) attack An attack in which many computers are hijacked and used to flood the target with so many false requests that the server cannot process them all, and normal traffic is blocked.

hacker Anyone who attempts to gain access to unauthorized resources on a network, usually by finding a way to circumvent passwords, firewalls, or other protective measures.

hactivists Computer attackers with political goals.

integrity The accuracy and consistency of information during its creation, transmission, and storage.

intrusion detection and prevention system (IDPS) A security tool used to detect and sometimes prevent an attack.

logic bomb Malware designed to be used at a specific time in the future or when a specified condition exists.

macro A type of script that automates repetitive tasks in Microsoft Word or similar applications.

mandatory access control (MAC) An access control method that defines an uncompromising manner for how information can be accessed. With the MAC method, all access capabilities are defined in advance.

nonrepudiation The capability to prevent one participant in an electronic transaction from denying that it performed an action.

packet filters Hardware or software tools that allow or deny packets based on specified criteria, such as port, IP address, or protocol.

packet monkeys A derogatory term for unskilled crackers or hackers who steal program code and use it in denial of service attacks instead of creating the programs themselves.

permissive policy A general approach to security that calls for a firewall and associated components to allow all traffic by default, blocking only specified traffic on a case-by-case basis.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

physical security Measures taken to physically protect a computer or other network device from theft, fire, or environmental disaster.

port An area in random access memory (RAM) reserved for the use of a program that "listens" for requests for the service it provides.

restrictive policy A general approach to security that calls for a firewall and associated components to deny all traffic by default, allowing only specified traffic on a case-by-case basis.

role-based access control (RBAC) An access control method that establishes organizational roles to control access to information. The RBAC method limits access by job function or job responsibility.

script kiddie A young, inexperienced computer programmer who spreads viruses and other malicious scripts and exploits weaknesses in computer systems using tools and techniques created by others.

scripts Executable code attached to e-mail messages or downloaded files that is used to infiltrate a system.

signature files Files used by antivirus programs that contain patterns of known viruses and malware.

signatures Signs of possible attacks that include an IP address, a port number, and the frequency of access attempts; an IDPS uses signatures to detect possible attacks.

socket A network connection consisting of a port number combined with a computer's IP address.

Trojan program A harmful computer program that appears to be something useful to deceive a user into installing it.

virtual private network (VPN) A network, typically the Internet, used to transmit confidential data secured by encryption, encapsulation, and authentication.

virus Computer code that copies itself from one place to another surreptitiously and performs actions that range from benign to harmful. Viruses require some user action, such as clicking an executable attachment or viewing an infected Web page, to enable them to launch.

virus scanning The process of examining files or messages for filenames, patterns, extensions, and other indications that a virus or other malware is present.

worm Computer files that copy themselves repeatedly and consume disk space or other resources. Worms do not require user intervention to be launched; they are self-propagating.

Review Questions

1. Most network threats originate from which of the following?
   a. inside the company
   b. script kiddies
   c. back doors
   d. industrial spies

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

2. What are some of the reasons for network attacks?
   a. industrial espionage
   b. revenge
   c. financial gain
   d. all of the above

3. The capability to prevent one participant in an electronic transaction from denying that it performed an action is called _____.
   a. plausible deniability
   b. integrity
   c. nonrepudiation
   d. undeniability

4. Servers with outside access to the public should be located on _____. (Choose all that apply.)
   a. their own subnet
   b. a DMZ
   c. an internal LAN
   d. a network perimeter

5. Packet filters can block or allow transmission of packets based on which of the following criteria? (Choose all that apply.)
   a. port number
   b. open ports
   c. time of access attempts
   d. IP address

6. An attacker who causes harm to systems in support of some principle is categorized as which of the following?
   a. cracker
   b. hacker
   c. industrial spy
   d. cyberterrorist

7. An IP address combined with a TCP/IP port number is called which of the following?
   a. network address
   b. socket
   c. script
   d. port ID

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

8. Firewall enforcement of policies is handled primarily through setting up packet-filtering rules, a set of which is contained in the _____.
   a. routing table
   b. rule base
   c. access control list
   d. packet filter

9. Name four goals of network security.

10. An uninterruptible power supply is a component of _____ security.
    a. virtual
    b. auditing
    c. physical
    d. password

11. The Stuxnet worm was designed to _____.
    a. shut down Internet DNS servers
    b. disrupt computer-controlled industrial operations
    c. steal financial information
    d. be used by script kiddies

12. A packet-filtering device evaluates data in the payload and compares it with a predefined set of rules. True or False?

13. Which of the following malware is designed to replicate itself? (Choose all that apply.)
    a. worm
    b. virus
    c. Trojan horse
    d. SYN flood

14. In a restrictive firewall policy, what is the starting point for developing a rule base?
    a. allow all traffic
    b. block all traffic except specified types
    c. allow all traffic except specified types
    d. block all traffic

15. In an IDPS, specific indications of a possible attack are called _____.
    a. signatures
    b. signals
    c. true positives
    d. alerts

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Projects

The hands-on projects in this book require one or more of the following operating systems: Windows 7 Professional, Windows Server 2008 R2 SP1, and Ubuntu desktop. Students can work in teams of two with one system running Windows Server 2008 as a real or virtual machine, and the other running Windows 7 and Ubuntu as a dual boot or as virtual machines. Recommended initial configurations are detailed in the Introduction of this book.

Hands-On Project 1-1: Installing a Personal Firewall

Time Required: 20 minutes

Objective: Install ZoneAlarm Free Firewall.

Description: Freeware firewall programs are not as full featured as commercial firewalls, but they are useful for testing and learning how the programs work. In this activity, you use a Windows 7 computer to download and use the freeware version of a popular personal firewall, ZoneAlarm by Check Point. Both Windows Server 2008 and Windows 7 are used in this project.

1. Boot the Windows Server 2008 system so that the Windows 7 user can log on to the domain. Log on to Windows 7 with an administrative account.
2. To disable the Windows firewall, click the **Start** button, and click **Control Panel**. If necessary, click **Category** in the upper-right corner and select **Small icons**. Click **Windows Firewall**, and click **Turn Windows Firewall on or off**. In each of the three sections—Domain network location settings, Home or work (private) network location settings, and Public network location settings—click the option button next to **Turn off Windows Firewall (not recommended)**. Click **OK** and close the Windows Firewall window.
3. Start your Web browser, and go to **http://www.zonealarm.com/security/en-us/zonealarm-pc-security-free-firewall.htm**.
4. Click the **Download** button to the right of ZoneAlarm Free Firewall to start the download. If the Internet Explorer bar appears and informs you that the file download has been blocked, click the bar and click **Download File**. Click **Save** and download the program to your Downloads folder.
5. When the download is finished, click **Open Folder**, right-click the downloaded program file, and click **Run as administrator**. If necessary, click **Yes** in response to the User Account Control warning.
6. Click the check box to accept the terms of the license agreement, and click **Next**.
7. Click to uncheck the top check box, which improves your Internet protection with the ZoneAlarm Security Toolbar. Click **Next**.

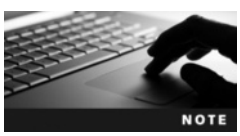

These options can be added any time after the installation is completed.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

8. Enter your name, uncheck the **Inform me about product updates and security news.** check box (see Figure 1-5), and click **Next**.

Source: ZoneAlarm

Figure 1-5 ZoneAlarm product registration

9. When the program setup has finished installing, click **Finish**.
10. If necessary, start ZoneAlarm by clicking Start and then clicking ZoneAlarm Security. When ZoneAlarm starts, it displays various windows. Close all of these windows except for the main window shown in Figure 1-6.

Source: ZoneAlarm

Figure 1-6 ZoneAlarm main window

11. Click **Update** to verify that you have the latest signatures. When the update is complete, click **Close**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

12. In the main window, click the **Computer Protected** box. The Computer tab that appears should look similar to Figure 1-7.

Source: ZoneAlarm

Figure 1-7 ZoneAlarm Computer tab

13. Click the **Basic Firewall** link. How many zones are listed? What are the functions of the zones? What zone settings are associated with medium and high security?
14. Click the **Advanced Settings** button. Notice the various settings; for example, VPN protocols are allowed by default.
15. Click **OK** to close the Firewall Settings window and return to the main Computer tab. What functions are not available in this free version? Click the **Application Control** link, investigate the settings that you can configure on the main Application Control page, and investigate the Advanced Settings.
16. Click the **Internet** tab. What functions are not available in this free version? What function is enabled on the Internet tab?
17. Click the **Identity & Data** tab and investigate the functions that are and are not available.
18. Leave the system running for the next Hands-On Project.

Hands-On Project 1-2: Installing a Port Scanning Tool

Time Required: 10 minutes

Objective: Install Nmap.

Description: One of the first steps in attacking a network, either by a network administrator who is testing for vulnerabilities or by a hacker looking for a system to exploit, is identifying the systems on the network. Learning a network's IP addresses, MAC addresses, operating

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

system, and services is a good first step for an attack. Port scanning tools like Nmap are used to get this basic information. In this activity, you install the Nmap tool.

1. Log on to Windows Server 2008 and Windows 7 with an administrative account.
2. On the Windows Server 2008 computer, open your Web browser and go to **http://nmap.org/download.html**.
3. Scroll down to the Microsoft Windows binaries section and click the link labeled as the Latest stable release self-installer. (See Figure 1-8.) If the Internet Explorer warning window appears, click **Add**. In the Trusted sites window, click to uncheck the box next to **Require server verification (https:) for all sites in this zone**. Click **Add**, click **Close**, and then click the self-installer link. If a window appears and asks if you want to run or save the file, click **Save**, and then save the file to your Downloads folder. When the download is complete, click **Run**. If a User Account Control window appears, click **Yes**.

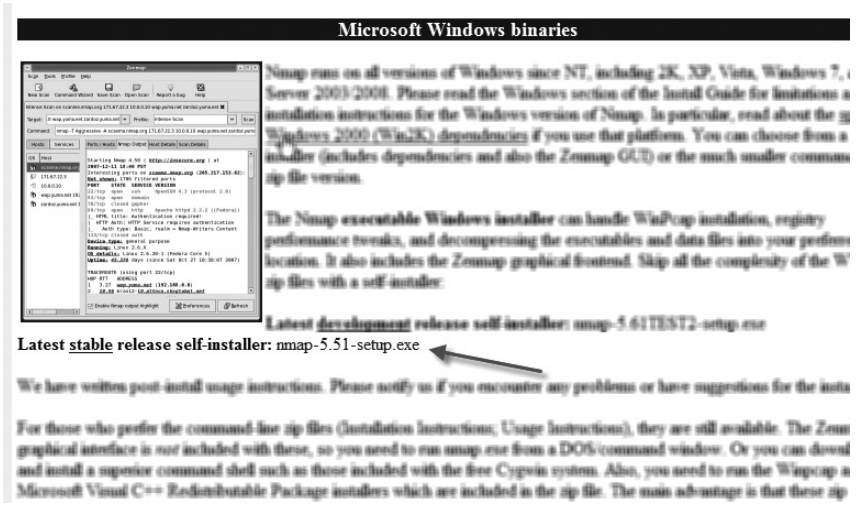

Source: Nmap

Figure 1-8 Nmap Windows binaries

4. In the License Agreement window, click **I Agree**. In the Choose Components window, click **Next**. In the Choose Install Location window, click **Install**.
5. During the installation of Nmap, you need to install the WinPcap program. In the WinPcap License Agreement window, click **I Agree**. In the Installation Complete window, click **Next**. In the WinPcap Options window, click **Next**. In the Finished window, click **Finish**. In the Installation Complete window, click **Next**. In the Create Shortcuts window, click **Next**. In the Finished window, click **Finish**.
6. Leave the system running for the next activity.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Project 1-3: Port Scanning a Target System

Time Required: 20 minutes

Objective: Obtain information about a target system using port scanning and identify the difference in results depending on firewall configuration.

Description: In this project, you scan Windows 7 under two conditions. First, you use the default firewall settings of ZoneAlarm. Because you will be scanning from a computer on the same network as Windows 7, you will be scanning from a "trusted" computer. Then you change the ZoneAlarm firewall settings to see how it responds when the attack comes from an "untrusted" computer.

1. Log on to Windows 7 with an administrative account.
2. Click the arrow on the right side of the Windows taskbar to show hidden icons. Verify that ZoneAlarm is running by checking the ZoneAlarm icon. (See Figure 1-9.)

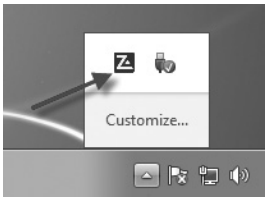

Source: ZoneAlarm

Figure 1-9 ZoneAlarm icon

3. Access the desktop of Windows Server 2008. Double-click the **Nmap** shortcut to start Nmap.
4. In the Target box, enter the IP address of the Windows 7 system. (You can find a system's IP address as follows: Click **Start**, click **All Programs**, click **Accessories**, click **Command Prompt**, enter **ipconfig**, and press **Enter**.) Verify that the Profile box reads "Intense scan." Click the **Scan** button. When the scan has finished, use the information obtained to complete the following table.

Total number of closed ports	
List TCP ports open	
List functions of TCP ports open	
NetBIOS name	
MAC address	
Operating system identified	

5. If the ZoneAlarm interface has been closed, go to the Windows 7 system, click the arrow that shows hidden icons in the taskbar, and click the **ZoneAlarm** icon. Click the **Computer Protected** box. Under Basic Firewall, note the number of access attempts blocked and record them here:

   _____

   Click **Basic Firewall**, and then use the slider bar to change the current settings of Your Trusted Zone from medium to high.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

6. From Windows Server 2008, access the Nmap program and verify that the Target box displays the IP address of the Windows 7 system. Verify that the Profile box reads "Intense scan." Click the **Scan** button. How does the Windows 7 system react? When the scan has finished, use the information obtained to complete the following table.

Total number of closed ports	
List TCP ports open	
List functions of TCP ports open	
NetBIOS name	
MAC address	
Operating system identified	

7. On the Windows 7 system, click the arrow that shows hidden icons in the taskbar, and click the **ZoneAlarm** icon. Click the **Computer Protected** box. Under Basic Firewall, note the number of access attempts blocked and record them here:

   _____

   How are these different from those noted in Step 5?

8. How are the scans from Steps 4 and 6 different? Given the number of security incidents that originate inside a trusted network, what is your assessment of ZoneAlarm's default configuration?
9. Leave Windows 7 running for the next activity.

Hands-On Project 1-4: Viewing ZoneAlarm Log Files

Time Required: 10 minutes

Objective: Examine a log file using a program designed for the task.

Description: In this project, you use a third-party program called Sawmill to help analyze a log file. Because the raw data in a log file can be confusing, log analysis software makes the process much easier and faster. To use the Sawmill software required in this project, a 64-bit operating system is recommended.

1. From the Windows 7 system, start your Web browser, and go to **http://www.sawmill.net/cgi-bin/download.pl**.
2. Click the option button next to **Do not send me email**, and click **Proceed**.
3. Click **Proceed**. Verify that Windows 7 is selected in the drop-down box, and click **Proceed**. In the next drop-down box, verify that your hardware platform (32-bit or 64-bit) is selected, and click **Proceed**. (If you are not sure which platform you are using, click **Start**, right-click **Computer**, click **Properties**, and verify that your platform is listed.)
4. If the Internet Explorer bar appears and informs you that the file download has been blocked, click the bar and then click **Download File**. Click **Retry** if the page needs to resubmit your information. In the window that appears, click **Save**, and

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

then save the file to your Downloads folder. If a User Account Control window appears, click **Yes**. When the download is complete, click **Run**. If necessary, click **Run** again in the Internet Security window.

5. In the Release Notes window, click **Next**. In the Choose Install Location window, click **Next**. In the Choose Components window, click **Install**. If ZoneAlarm displays a warning that Sawmill is trying to use Windows Installer, click **Allow**. In the final wizard screen, click **Finish**.
6. In the Welcome to Sawmill window, click **Next**. In the License Agreement window, click the option button to accept the terms in the license agreement, and click **Next**. In the Licensing window, verify that the option to try Sawmill for 30 days is selected, and click **Next**. In the Root Administrator window, enter **Administrator** in the Username box; in the Password and Reenter password boxes, enter **Pa$$word** and then click **Next**. In the Email and SMTP Server window, click **Next**. In the Trial Edition window, click the **Enterprise** option button and click **Next**. In the Automated Feedback Agent window, uncheck the box that allows Sawmill to send you information, and click **Next**. In the Antivirus Warning window, click **Next**. In the Complete Setup window, click **Finish**.
7. In the Create new Profile window, click **Start here**. In the Log source window, click the Log source drop-down box to examine the possible sources of log files, but select **Local disk or mapped/mounted disk**. You will use the ZoneAlarm log file location. Open ZoneAlarm, click the **Tools** link, click **Logs**, click **Log Control**, and note the path to the Log Archive Location. Enter this path in the Pathname box in the Sawmill Log source window, and click **Next**. In the Log format detected window, click **Next**. In the Database window, click **Next**. In the Database performance options window, click **Next**. In the Profile name window, enter **ZoneAlarm** in the Profile name box, and click **Finish**. In the next window, click the **Process Data & View Reports** link.

   The Sawmill Overview window should open. It should be similar to Figure 1-10.

   Note the different views listed on the left side of the window. Explore views such as Date/times and Hours of day. Because you have not logged activity over an extended time, these views are not very revealing, but the ability to zero in on events by days of the week and hours of the day enables administrators to construct patterns that can help them anticipate security incidents.
8. Click the **Types** link to see the types of events that were logged. See Figure 1-11.
9. Click the **Source IPs** link. Either all or most events will show an empty IP address. These are the events generated by Nmap. In Figure 1-12, one public IP address from the Internet was captured.

   Figure 1-13 shows the result of clicking the **Cities** link. You can see that by using Internet resources, Sawmill has determined that the event's source was Cambridge, Massachusetts. Spend some time to explore the other view links on the left side of the window.
10. Return to the ZoneAlarm window. Click **Tools**, click **Logs**, click **Log Control**, and then click **View Log** to see the log file that Sawmill analyzed. Analyzing this text file can be quite tedious, but a graphical tool like Sawmill organizes logged information into easy-to-read columns and lets you sort them by date, IP address, or other criteria.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

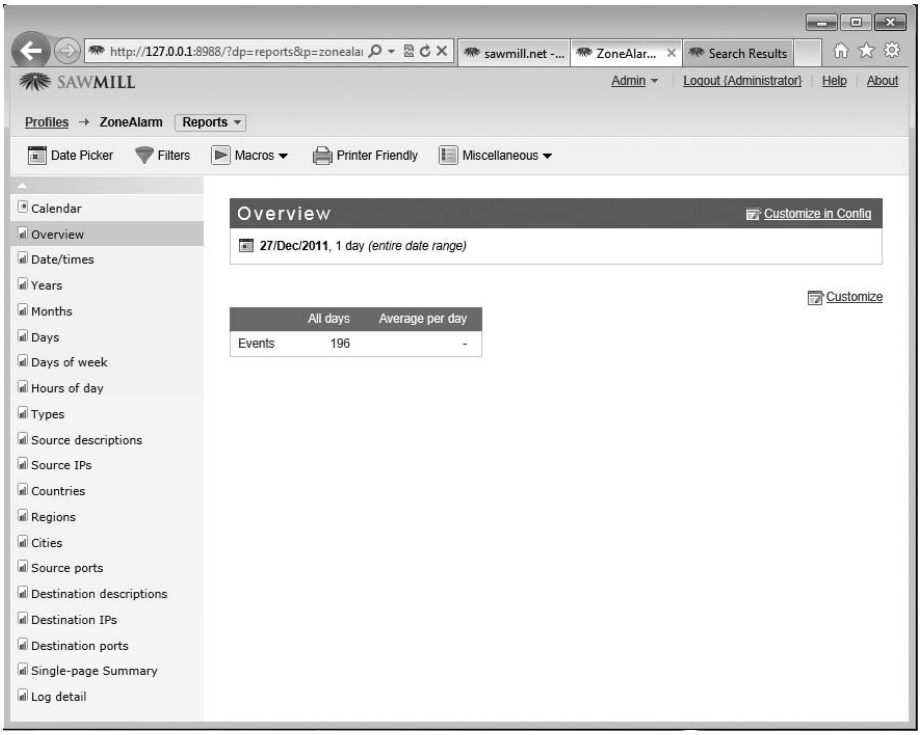

Source: Sawmill

Figure 1-10 Sawmill overview

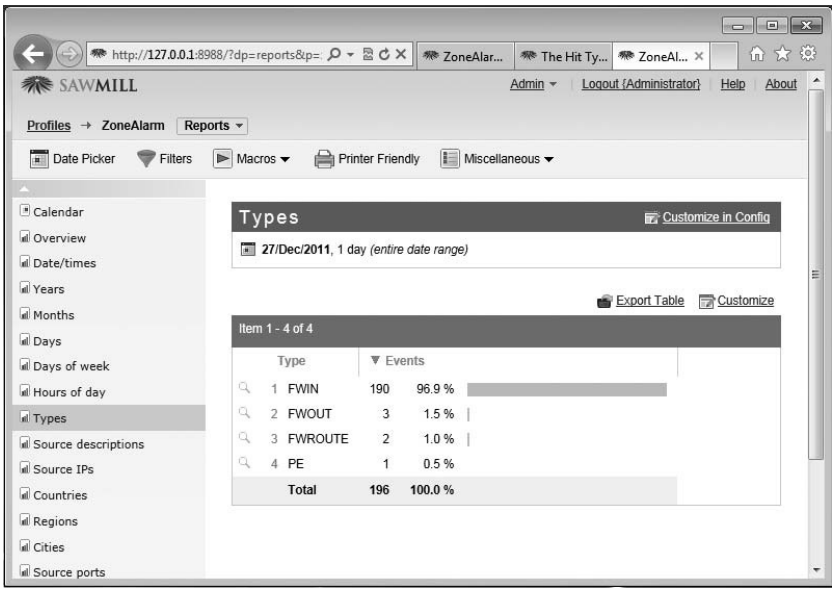

Source: Sawmill

Figure 1-11 Types of logged events

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

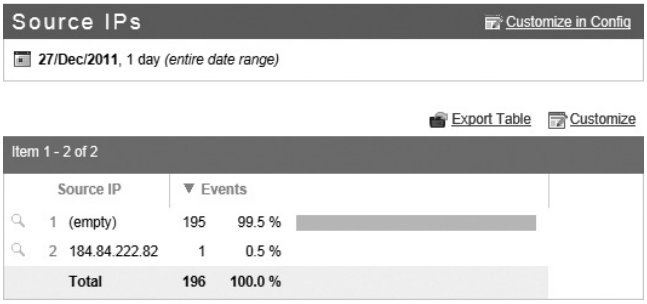

Source: Sawmill

Figure 1-12 Source IP addresses of logged events

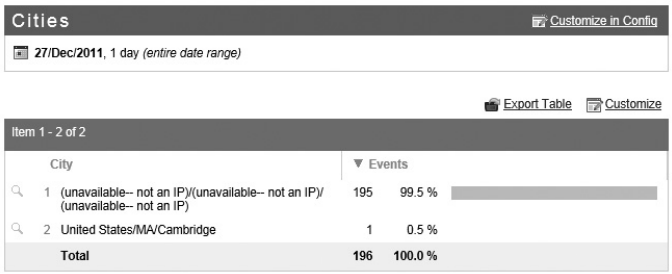

Source: Sawmill

Figure 1-13 Source cities of logged events

11. From the ZoneAlarm window, click the **Computer Protected** link, click **Basic Firewall**, and set Your Trusted Zone security to **Medium**.
12. Log out of Windows 7. If necessary, log out of Server 2008.

Case Projects

Case Project 1-1: Determining Legal Requirements for Penetration Testing

Alexander Rocco Corporation, a large real estate management company in Maui, Hawaii, has contracted your computer consulting company to perform a penetration test on its computer network. Penetration testers are hired to attack an organization's network, determine what vulnerabilities are present, and provide recommendations for securing the company's information systems. The management company owns property that houses a five-star hotel, golf courses, tennis courts, and restaurants. Claudia Mae, the vice president, is your only contact at the company. To avoid undermining your tests, you will not be introduced to any IT staff or employees. Claudia wants to determine what you can find out about the company's network infrastructure, network topology, and vulnerabilities without any assistance from her or company personnel.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Based on this information, write a report outlining the steps you should take before beginning the penetration tests. Research applicable state and federal laws, and reference them in your report.

Case Project 1-2: Understanding the Rules of Engagement for Security Testers

You are a new security tester for Security Consulting Company (SCC). Before you begin your first assignment, Shelley Canon, the vice president of SCC, wants you to read the rules of engagement section of the Open Source Security Testing Methodology Manual (OSSTMM) to make sure you do not violate any company policies.

Write a memo to Shelley that summarizes the OSSTMM's rules of engagement section. The memo should describe the purpose of the rules of engagement and include answers to the following questions:

- When is it permissible to release the names of past clients?
- If you are not able to penetrate a client's network, is it permissible to offer your services free of charge?
- When is it permissible to conduct denial of service attacks on a client's network?

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 2

TCP/IP

After reading this chapter and completing the exercises, you will be able to:

- Explain the fundamentals of TCP/IP networking
- Describe IPv4 packet structure and explain packet fragmentation
- Describe Internet Protocol version 6 (IPv6)

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

This chapter introduces you to the fundamentals of TCP/IP networking. If you want to secure a network, understanding the TCP/IP protocol suite is vital. You review IP addressing briefly, revisit the TCP life cycle, and review how TCP communications sessions work. Then you move on to learn about Internet Protocol version 6 (IPv6), the next generation of IP.

Even if you are familiar with TCP/IP and IPv4, IPv6 involves some new protocols and functions you need to know. In this chapter, you examine the core protocols of IPv6 and learn about IPv6 addressing and utilities.

The OSI Model and TCP/IP Protocols

Transmission Control Protocol/Internet Protocol (**TCP/IP**) is a suite of many protocols for transmitting information from point to point on a network. (This suite of protocols is often called a stack.) You may already know TCP/IP, but it never hurts to brush up on your knowledge. Given the dynamic nature of information technology, you can forget even simple details, so this section gives you a quick refresher on TCP/IP basics, such as the Open Systems Interconnection (OSI) model, IP addressing, and subnetting.

The OSI Model

You are probably familiar with the OSI reference model of network communications, which divides the communication functions used by two hosts into seven separate layers. TCP/IP has its own stack of protocols that correspond roughly to these layers. Table 2-1 compares the two models.

OSI	TCP/IP stack
Application	HTTP DNS DHCP
Presentation	FTP SNMP Telnet
Session	IMAP SMTP POP
Transport	TCP UDP
Network	IP ICMP IGMP ARP RIP OSPF
Data Link	Device Drivers
Physical	Network Adapter

© Cengage Learning 2014

Table 2-1 The OSI model and the subprotocols of the TCP/IP stack

You should be familiar with most of these protocols and their functions. The following list provides a brief review of some protocols; more detailed discussions of TCP, UDP, IP, ICMP, IGMP, and ARP are included later in this chapter. If you need a more detailed refresher on the TCP/IP stack, the OSI model, and the major protocols operating at different layers, perform an Internet search on "TCP/IP and the OSI model." Dozens of helpful sites, such as *www.tcpipguide.com/free/index.htm*, are available for every level of knowledge.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Note that the top three layers of the OSI model—Application, Presentation, and Session—are considered equivalent to the top layer of the TCP/IP stack. The TCP/IP subprotocols listed in this layer are services that support a number of network functions:

- HTTP (Hypertext Transfer Protocol) is responsible for the delivery of Web documents formatted in HTML (Hypertext Markup Language) and other similar languages.
- DNS (Domain Name System) is responsible for the resolution of fully qualified domain names (for example, support.microsoft.com) to IP addresses, as well as resolution of IP addresses to fully qualified domain names.
- DHCP (Dynamic Host Configuration Protocol) is responsible for automatic assignment of IP addresses and other configuration data to client systems.
- FTP (File Transport Protocol) provides efficient delivery of files from one system to another.
- SNMP (Simple Network Management Protocol) monitors various parameters on network devices like switches and servers.
- Telnet is a terminal emulation service that allows users to run commands on a remote system.
- IMAP, SMTP, and POP are used in e-mail communications. Internet Message Access Protocol is used for downloading e-mail, Simple Mail Transfer Protocol is used for sending e-mail, and Post Office Protocol, like IMAP, downloads e-mail. Unlike IMAP, however, POP does not allow the user to determine which messages are removed from the e-mail server.

The TCP/IP routing protocols RIP and OSPF are processed at the Network layer of the OSI model. These protocols allow routers to share their routing tables with each other. (Routing is covered in detail in Chapter 4.)

TCP/IP Addressing

IP addresses are one of the methods used to identify computers. These addresses are processed at the Network layer of the OSI model. The type of IP addresses most commonly in use conform to **Internet Protocol version 4 (IPv4)**, which specifies addresses with 32 bits of data. Each 32-bit address is divided into four groups called octets; each octet contains 8 bits of data. In binary, an IP address looks like this:

10000000.00100110.00101100.11100010

While binary numbers are no problem for computers, this notation is difficult for humans to manage; therefore, IP addresses are usually converted to dotted decimal notation, such as 192.168.10.1. An IP address consists of two main parts:

- The **network identifier**, which is the part of an IP address shared among computers in a network segment
- The **host identifier**, which is unique to each computer on the network segment

These two identifiers are defined by another dotted decimal value called the **subnet mask**. This value indicates which part of the IP address is the network identifier and which part is the host identifier. You learn more about subnet masks in the "Subnetting" section later in this chapter.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

One way attackers can gain access to your network is by determining the IP addresses of computers. After they have an address, they can attempt to take over the computer and use it to launch attacks on other computers in the network or access network resources. Therefore, a fundamental requirement of network security is to understand IP addresses and other network addresses so that you can conceal or change them to deter attackers.

IP addresses are valuable commodities. If attackers can find a computer's IP address, they can run a port scan to look for open ports they can exploit. By hiding IP addresses, you can prevent certain attacks. To hide the addresses of computers on your network, you can use **Network Address Translation** (**NAT**) to translate your private network's internal addresses into the address of the NAT server's external interface connected to the Internet. A private network's internal addresses are not routable on the Internet.

Security is not the only reason for using NAT. The Internet has grown more quickly than expected by the creators of the IPv4 32-bit addressing scheme. Today, IP addresses are in short supply, so Internet Protocol version 6 (IPv6) is being implemented.

Address Classes

IPv4 addresses are separated into address categories called classes. An IP address class is determined by the number of its networks compared to the number of its hosts. For example, a Class A address uses 8 bits for the network portion of the address and 24 bits for the host portion. The class divisions are shown in Table 2-2.

Class	First octet decimal range	Default subnet mask	Purpose
Class A	1–126 127.x.x.x is reserved; the address 127.0.0.1 is used to indicate the local system's TCP/IP implementation	255.0.0.0	Large corporations and governments
Class B	128–191	255.255.0.0	Medium networks
Class C	192–223	255.255.255.0	Small networks
Class D	224–239	N/A	Multicasting
Class E	240–254	N/A	Experimentation

© Cengage Learning 2014

Table 2-2 IP address classes

Private IP Address Ranges

In addition to public address ranges, which are used to move messages from one network to another, protocol designers recognized the need for a private addressing system that organizations could use to build internal infrastructures. To obtain a public IP address, individuals and organizations must register and pay a fee for each address. Addresses with large quantities of hosts carry a higher price tag than those with fewer host addresses. The private addressing scheme eliminated the need to purchase addresses for every group of machines. However, the free address space carries a price. Private addresses are not routable over the Internet. RFC (Request for Comments) 1918 defined ranges of reserved private IP addresses

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

that organizations can use on their internal networks. No Internet host can directly access an organization's computers as long as they have only private IP addresses. Table 2-3 lists the private IP address ranges.

RFCs are the documents that create standards for Internet technology. You should become familiar with RFCs so that you can keep up with changes as they occur. You can read RFC 1918 at *http://tools.ietf.org/html/rfc1918*.

Network address	Subnet mask	First valid host address	Last valid host address	Broadcast address
10.0.0.0	255.0.0.0	10.0.0.1	10.255.255.254	10.255.255.255
172.16.0.0	255.240.0.0	172.16.0.1	172.31.255.254	172.31.255.255
192.168.0.0	255.255.0.0	192.168.1.1	192.168.255.254	192.168.255.255

© Cengage Learning 2014

Table 2-3 Private IP address ranges

Subnetting

As mentioned, an IP address and its subnet mask identify both the network and the host. Address classes already have network identification octets in the subnet mask set by default; Class A has the first octet set, Class B the first two, and Class C the first three (see Table 2-2). The remainder of the address is available for extending the network identifier portion of the subnet mask. For example, a default Class B address has 16 bits available for the host portion of the address. This means that a single Class B network has more than 65,000 host addresses! However, many organizations use some of the host bits to identify the network; this approach creates several smaller subnetworks that are more flexible and easier to manage than one large network.

Subnetting is used to segment internal networks logically, as described previously. It can also be used for the following purposes:

- Mirroring the organization's physical layout
- Mirroring the organization's administrative structure
- Planning for future growth
- Reducing and controlling network traffic
- Increasing network security

When administrators create subnetworks that mirror the organization's structure or physical layout, managing security, access needs, and auditing becomes easier. This concept is similar to creating sites and domains. Trying to manage many different user groups gets complicated. If all users with similar security and access needs are grouped into a single subnet, the entire group can be managed instead of managing each user separately. Subnets are used to make network management easier and to optimize security, performance, and access.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

When administrators create subnets, they borrow bits from the host portion of the IP address to make a set of subnetworks. The number of borrowed bits determines how many subnets and hosts are available. An administrator can use up to 14 bits for subnetting a Class B network. (Two bits must be available for hosts.) When you subnet, you lose some addresses that would have been available for hosts, but a network of 65,000 hosts would be unwieldy to manage, and performance would suffer with so much traffic on a single network segment. Table 2-4 shows how subnetting is used on Class B networks.

Subnet	Number of subnetworks	Usable hosts per subnet
255.255.128.0	2	32766
255.255.192.0	4	16384
255.255.224.0	8	8190
255.255.240.0	16	4094
255.255.248.0	32	2046
255.255.252.0	64	1022
255.255.254.0	128	510
255.255.255.0	256	254
255.255.255.128	512	126
255.255.255.192	1024	62
255.255.255.224	2048	30
255.255.255.240	4096	14
255.255.255.248	8192	6
255.255.255.252	16384	2

© Cengage Learning 2014

Table 2-4 Class B subnetting

Although you should already have a good working knowledge of subnetting, a brief refresher is useful. You might find it helpful to remember how the binary numbering system stacks up to the decimal system in IP addressing. Table 2-5 shows the decimal equivalents of binary place values.

Binary digit	1	1	1	1	1	1	1	1
Decimal equivalent	128	64	32	16	8	4	2	1

© Cengage Learning 2014

Table 2-5 Binary-to-decimal values

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Now that you remember what decimal value is represented by each binary digit in an IP octet, you can review subnetting. First, you select the mask that meets your needs, and then you assign the mask to your network. To do this, you must find the usable IP address ranges for the network address, the valid host address range, and the broadcast address. Converting the last masking octet to binary performs this task most easily. The binary place value of the last masking digit is the block size. (Block size refers to the maximum number of host addresses plus the subnetwork and broadcast address in a subnet.) Subnetting a Class C address goes like this:

Class C address: 199.1.10.0 (network address)

Default mask: 255.255.255.0

Selected mask: 255.255.255.224

Mask converted to binary: 11111111.11111111.11111111.11100000

Notice that the last masked digit occupies the binary place value of 32 (which is the block size). Starting with the network address, increment by 32 until you reach the mask's number (224). It looks like this:

- 0 (00000000)
- 32 (00100000)
- 64 (01000000)
- 96 (01100000)
- 128 (10000000)
- 160 (10100000)
- 192 (11000000)
- 224 (11100000)

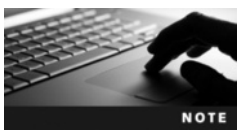

If you first learned to subnet some years ago, you might have learned that the subnetwork identifier cannot be all binary zeros or all binary ones. This was true when Routing Information Protocol version 1 (RIPv1) was in use. With subsequent routing protocols such as RIPv2 and OSPF (Open Shortest Path First), all zeros and all ones are permitted in the subnet identifier.

Table 2-6 shows the subnetwork addresses, host address ranges, and broadcast addresses for this example.

When looking at the chart, you might notice that the host address range includes only 30 valid hosts. What happened to the other two hosts? When you calculate a host range, or block size, you include one address for the network address and one for the broadcast address. These two addresses cannot be assigned to hosts. Therefore, you can use the following formula to calculate the number of valid hosts:

$2^x - 2$ = Number of valid hosts per subnet (the exponent x is the number of bits in the host identifier)

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Subnet address	Valid host address range	Broadcast address for subnet
199.1.10.0	199.1.10.1-199.1.10.30	199.1.10.31
199.1.10.32	199.1.10.33-199.1.10.62	199.1.10.63
199.1.10.64	199.1.10.65-199.1.10.94	199.1.10.95
199.1.10.96	199.1.10.97-199.1.10.126	199.1.10.127
199.1.10.128	199.1.10.129-199.1.10.158	199.1.10.159
199.1.10.160	199.1.10.161-199.1.10.190	199.1.10.191
199.1.10.192	199.1.10.193-199.1.10.222	199.1.10.223
199.1.10.224	199.1.10.225-199.1.10.254	199.1.10.255

© Cengage Learning 2014

Table 2-6 Subnetting example

Variable Length Subnet Masking

Networks that do not have a large number of available IP addresses can use **variable length subnet masking** (**VLSM**), which involves applying masks of varying sizes to the same network. If an organization has a limited number of IP addresses and subnets of varying lengths, VLSM can help it use address space more efficiently. VLSM is a means of allocating IP addressing according to the network's needs. This allocation method creates subnets within subnets and multiple divisions of an IP network. VLSM is often used to secure stub networks—"dead end" networks that have only one connection to any other network. VLSM is also used to secure serial lines, which are connections between remote networks that require only two IP addresses. In these cases, VLSM makes the subnets only as large as needed.

Classless Interdomain Routing

Classless Interdomain Routing (**CIDR**) is an address notation scheme that specifies the number of masked bits in an IP address/subnet mask combination. Instead of using standard notation for subnet masks, with CIDR you can simply list the number of masked binary bits. The subnet mask 255.255.255.224, for example, has a total of 27 masked bits (eight in each of the first three octets and three in the last octet). In CIDR notation, you would write the network address 192.168.6.0 with a subnet mask of 255.255.255.224 as 192.168.6.0/27. CIDR overcomes the limitations of the default subnet masks of 8, 16, and 24 bits for Classes A, B, and C (classful addressing) so that unused addresses do not go to waste.

If you use subnet masks to segment network traffic into a series of smaller subnetworks, plan in advance how you will allocate nodes to each segment and assign subnet masks to those segments. Your planning might include network growth projections for the next two to five years so that you will not have to set up different subnet designations with each future network change.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Supernetting, also known as summarization, is used to summarize multiple routing table entries into one entry. In addition, classless routing is used to exchange subnet mask information between routers in routing updates. Classless routing allows VLSM and supernetting to work.

Unicasting, Multicasting, and Broadcasting

Each IP address class (Class A through Class D) is used with a different type of network. The address classes reflect the network's size and whether the packet is unicast or multicast. In a **unicast** transmission, one packet is sent from a server computer to each client computer that requests a file or an application, such as a streaming video presentation. If five clients request the video presentation, the server transmits the presentation separately to each client. In the same example, a **multicast** transmission means the server can treat all five clients as a group and send one transmission that reaches all of them. Multicasts can be used to reduce network traffic when transmitting bandwidth-intensive applications or files to multiple hosts. Instead of sending one to each host separately, the files or applications can be sent to all recipients at once.

A third type of network communication called a **broadcast** sends a communication to all points on a specific network. (Routers are usually configured so that they do not forward broadcasts to other networks.) There are two types of broadcasts:

- Flooded broadcasts are sent to any subnet. Routers do not forward the broadcasts because they are considered local. These broadcasts use the address 255.255.255.255.
- Directed broadcasts are sent to a specific subnet. Routers forward directed broadcasts using the broadcast address for the intended subnet.

Examining Internet Protocol Version 4 (IPv4)

TCP/IP is packet-based; it gives computers a fairly simple framework for transmitting information in small chunks called packets. Unfortunately, TCP/IP packets give attackers another way to gain entry into a network. Attackers can intercept packets and falsify the information in them or manipulate the packets in a way that makes it impossible for receiving servers to respond, which then disables those servers and opens the network to attack.

IP Datagrams

The portion of the packet that IP is responsible for routing through networks is called an IP **datagram**. This portion of the packet is processed at the Network layer of the OSI model. Each complete message is usually separated into multiple datagrams. Each datagram contains information about the source and destination IP addresses, a variety of control settings, and the actual data message that the computers are exchanging.

Each IP datagram is divided into different sections. The primary subdivisions are the header and the data, as described in the following sections. Some packets have another segmented section at the end called a **footer** (or "trailer"), which contains data that indicates the end of the packet. An error-checking algorithm called a Cyclic Redundancy Check (CRC) might also be added in the footer.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

IP Header Structure

The data in an IP packet is the part that end users see, but the **header** is the part that computers use to communicate, and it plays an important role in network security and intrusion detection. An IP header contains a number of fields and is similar to a TCP header, as you will learn in the "TCP Headers" section later in this chapter. Figure 2-1 shows a common way of depicting information in an IP header, which is divided into sections of 32-bit layers.

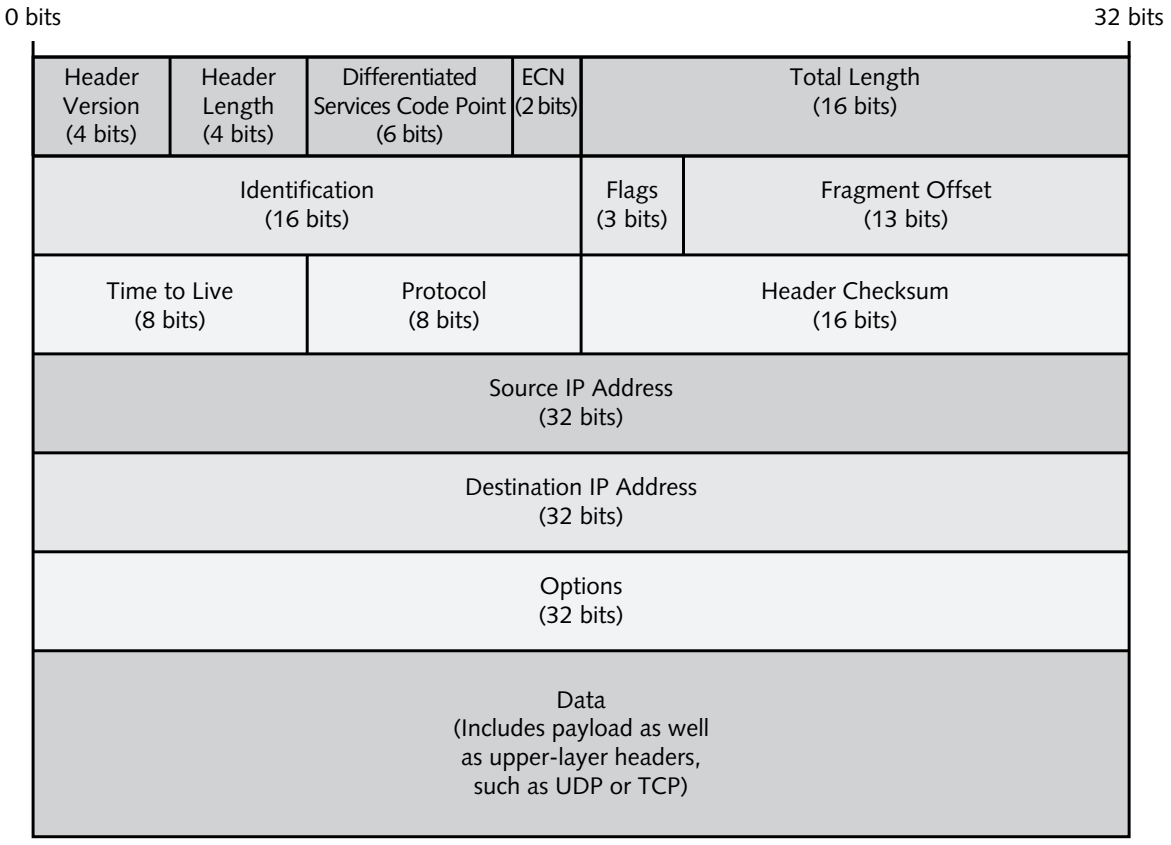

© Cengage Learning 2014

Figure 2-1 IP header structure

It is important to understand the different fields in an IP header so you can configure packet filters properly. Each field has varying importance to attackers, so you must know what each one does to protect against different types of attacks. The following fields make up the IP header structure:

- *Header Version*—This 4-bit field identifies the IP version used to generate the datagram. Because this is an IPv4 header, the value should indicate version 4.
- *Header Length*—This field describes the length of the header in 32-bit words, and is a 4-bit value. The default value is 20 bytes. Thus, a header length that has a binary size of 0101, which equals 5 in decimal, indicates 5 × 32 bits or 160 bits. Next, 160 bits/8 = 20 bytes. A "word" is the width of a computer processor's registers. As you should recall from your

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

hardware classes, registers are the storage areas where the processor does calculations. Most computers since the 386 processor have had a 32-bit word size, although most new computers now support a 64-bit word size and processors like the AMD Athlon or the Intel Core i7.

- *Differentiated Services Code Point (DSCP)*—This 6-bit field expresses the quality of service in the datagram's transmission through the network. This field is used by services that are sensitive to latency issues, such as Voice over IP (VoIP).
- *Explicit Congestion Notification (ECN)*—This 2-bit field allows ECN-compliant routers that are operating on ECN-compliant network infrastructures to signal congestion and thus minimize dropping of packets.
- *Total Length*—This 16-bit field specifies the datagram's total length to a maximum of 65,535 bytes.
- *Identification*—This 16-bit value helps divide the data stream into packets of information. The receiving computer (possibly a firewall) uses each packet's identification number to reassemble the packets that make up the data stream in the correct order.
- *Flags*—This 3-bit value indicates whether the datagram is a fragment—one datagram within a sequence of datagrams that make up an entire communication—and whether it is the last fragment or more will follow.
- *Fragment Offset*—If data is received in the form of a fragment, this value indicates where the fragment belongs in the sequence so that a packet can be reassembled.
- *Time to Live (TTL)*—This 8-bit value identifies the maximum amount of time the packet can remain in a network before it is dropped. Each router or device through which the packet passes (hops) reduces the TTL by a value of one. The TTL avoids congestion that results from corrupted packets infinitely looping through the network.
- *Protocol*—This field identifies the type of protocol being carried. For example, 1 = ICMP, 2 = IGMP, 6 = TCP, 17 = UDP, 47 = GRE (Generic Routing Encapsulation), 50 = ESP (Encapsulating Security Payload), and 51 = AH (Authentication Header).
- *Header Checksum*—This field is the sum of the 16-bit values in the datagram header; it is calculated at every hop to ensure accuracy of the header.
- *Source IP Address*—This field is the address of the computer or device that sent the IP datagram.
- *Destination IP Address*—This field is the address of the computer or device that received the IP datagram.
- *Options*—This field can include items such as a security field and several source routing fields that the packet sender uses to supply routing information. Gateways can then use this routing information to send the packet to its destination.

Programs that capture packets as they pass through a network interface give you another way to view packet header information. Most network operating systems (NOSs) have some type of built-in or add-on program to monitor network activity, such as Windows Network Monitor. Many security administrators, however, prefer third-party applications for their versatility and extra features. For example, Wireshark, which was formerly Ethereal, is an open-source network analysis utility that tracks packets and supplies detailed information

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

about them. (For more information, see *www.wireshark.org* and Hands-On Project 2-2 at the end of this chapter.) Figure 2-2 shows a portion of a Wireshark capture of a TCP/IP packet. The "A" indicates the beginning of the IP header. The fields are indicated by an "X." The "B" shows the start of the TCP header, the upper-level protocol that IP was transporting in this example.

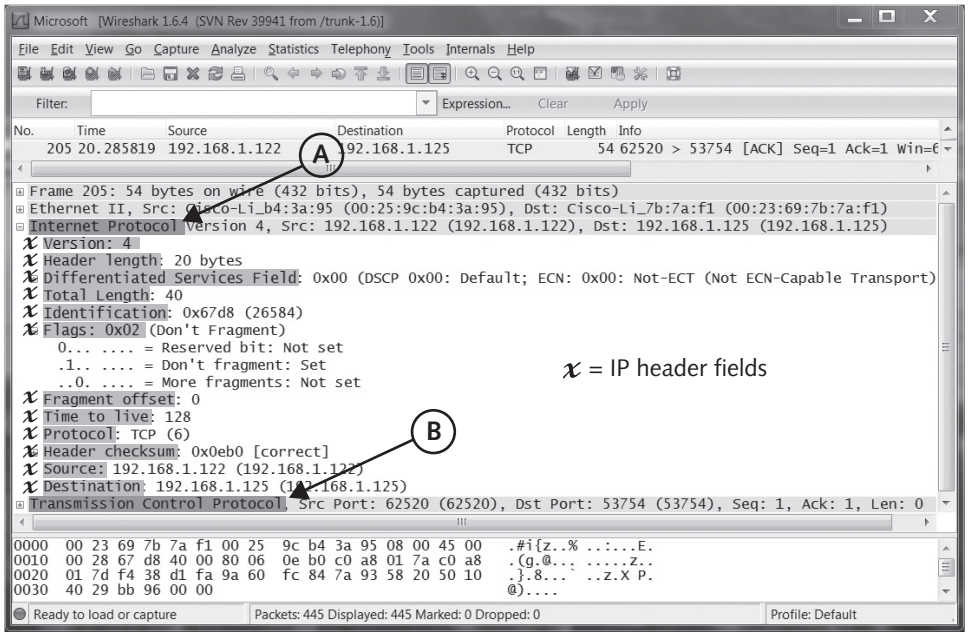

Source: Wireshark

Figure 2-2 IP header structure as seen in a Wireshark packet capture

ICMP Messages

Internet Control Message Protocol (ICMP) is designed to assist TCP/IP networks with troubleshooting communication problems. For example, using the ping command, ICMP produces messages that indicate whether a host has connectivity with another host. Like IP, ICMP is processed at the Network layer of the OSI model.

Based on a packet's message type, a firewall or packet filter must be able to determine whether an ICMP packet should be allowed to pass. An administrator might want to allow specific ICMP packet types for network diagnostics while wanting to block other types that could be used as part of an attack. Table 2-7 lists some common ICMP types.

Many ICMP types have codes associated with them. Some common codes are used for type 3 (Destination Unreachable) ICMP messages. For example, a type 3 ICMP message with a code of 13 indicates that the message was administratively prohibited, which usually means that an access list or firewall rejected the message. You can find a complete list of ICMP types and their codes at *www.iana.org/assignments/icmp-parameters/icmp-parameters.xml*.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

ICMP type	Name	ICMP type	Name
0	Echo Reply	17	Address Mask Request
3	Destination Unreachable	18	Address Mask Reply
4	Source Quench	30	Traceroute
5	Redirect	31	Datagram Conversion Error
6	Alternate Host Address	32	Mobile Host Redirect
8	Echo	33	IPv6 Where-Are-You
9	Router Advertisement	34	IPv6 I-Am-Here
10	Router Selection	35	Mobile Registration Request
11	Time Exceeded	36	Mobile Registration Reply
12	Parameter Problem	37	Domain Name Request
13	Timestamp	38	Domain Name Reply
14	Timestamp Reply	39	SKIP
15	Information Request	40	Photuris
16	Information Reply	1–2, 7, 19–29, 41–252	Unassigned or Reserved

© Cengage Learning 2014

Table 2-7 ICMP types

TCP Headers

TCP/IP packets do not contain just IP header information. They might also contain TCP headers (shown in Figure 2-3) that provide hosts with a different set of flags—and that give attackers a different set of components they can misuse in an attempt to attack networks. TCP headers are processed at the Transport layer of the OSI model. The TCP portion of a packet is called a TCP segment.

The Flags section of a TCP header is the set of nine 1-bit fields identified in Figure 2-3. From a security standpoint, the flags are important because you can specify them when you create packet-filtering rules. For example, the TCP header portion of a TCP packet that has an acknowledgement (ACK) flag set to 1 rather than 0 indicates that the destination computer received the packets that were sent. The first three flags—NS, CWR, and ECE—are related to Explicit Congestion Notification (ECN). The next six flags—URG, ACK, PSH, RST, SYN, and FIN—are of particular importance from a security perspective.

- *NS (Nonce Sum)*—Associated with ECN
- *CWR (Congestion Window Reduced)*—Associated with ECN
- *ECE (ECN Echo)*—Associated with ECN
- *URG (Urgent)*—When set to 1, data should be considered significant
- *ACK (Acknowledgement)*—Indicates that the previous packet was received

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

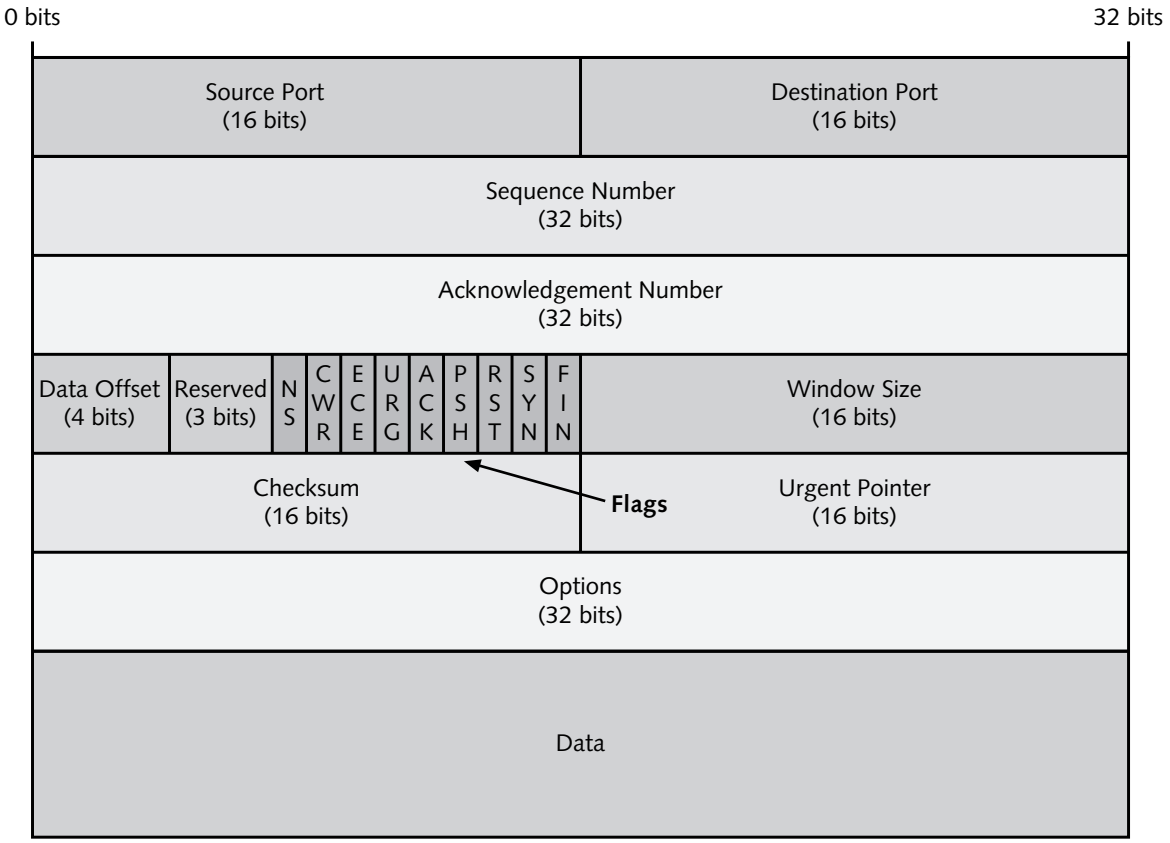

© Cengage Learning 2014

Figure 2-3 TCP header structure

- *PSH (Push)*—Forces TCP to deliver data rather than buffer it on the receiver
- *RST (Reset)*—Resets the connection
- *SYN (Synchronize)*—Synchronizes the sequence numbers
- *FIN (Finish)*—Indicates that no more data will come from the sender

Figure 2-4 shows how a TCP packet is displayed in a packet analyzer.

UDP Headers

The **User Datagram Protocol (UDP)**, like TCP, is processed at the Transport layer of the OSI model. This portion of a packet is called a UDP datagram and it provides a transport service for IP, but this protocol is considered unreliable because it is connectionless. In other words, a UDP packet does not contain the sequence number/acknowledgement number mechanism that enables TCP to guarantee delivery of the packet. (This connection-oriented mechanism is explained in detail later.) UDP is much faster than TCP because of the relative lack of overhead information in the header, and is appropriate when delivery does not need to be guaranteed. UDP simply sends the packets and relies on other protocols to ensure delivery,

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

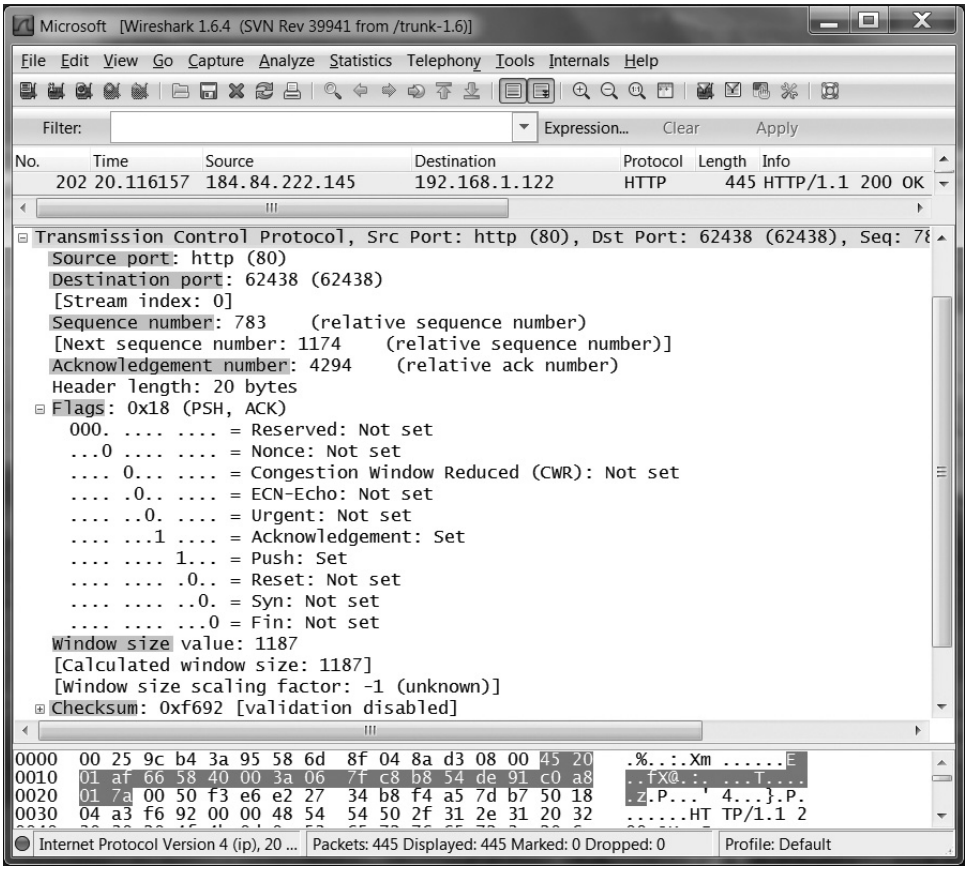

Source: Wireshark

Figure 2-4 TCP header structure as seen in a Wireshark packet capture

perform error checking, and so on. It does not provide errors and depends on the application in use to notify the user of any errors in transmission. UDP is especially useful for real-time applications, multimedia, or anything that requires speed over reliability.

UDP is used for broadcasting messages or for protocols that do not require the same level of service as TCP. For example, Simple Network Management Protocol (SNMP) and Trivial File Transfer Protocol (TFTP) are normally used on LANs, where packet loss is not considered a serious problem. On the other hand, attackers can scan for open UDP services to exploit by sending empty UDP datagrams to a suspected open port. If the port is closed, the system sends back an ICMP Destination Unreachable message (type 3). UDP packets have their own headers, as shown in Figure 2-5. As you can see, the UDP header is much smaller and simpler than a TCP header. In Figure 2-6, you can see how a UDP packet is displayed in a packet analyzer.

Packet Fragmentation

Fragmentation of IP packets was originally developed as a means of allowing large packets to pass through routers that had frame size limitations. Routers were then able to divide packets

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

into multiple fragments and send them along the network, where receiving routers reassembled them in the correct order and passed them along to their destination.

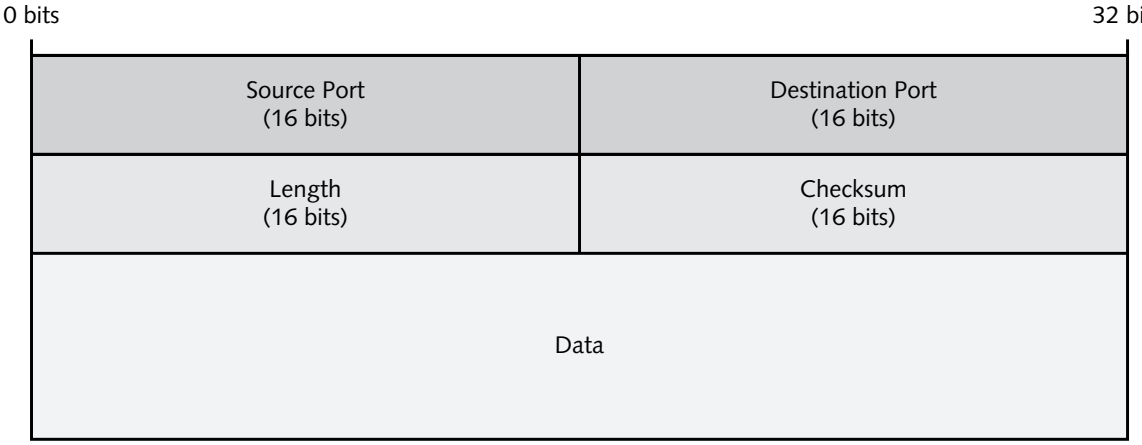

© Cengage Learning 2014

Figure 2-5 UDP header structure

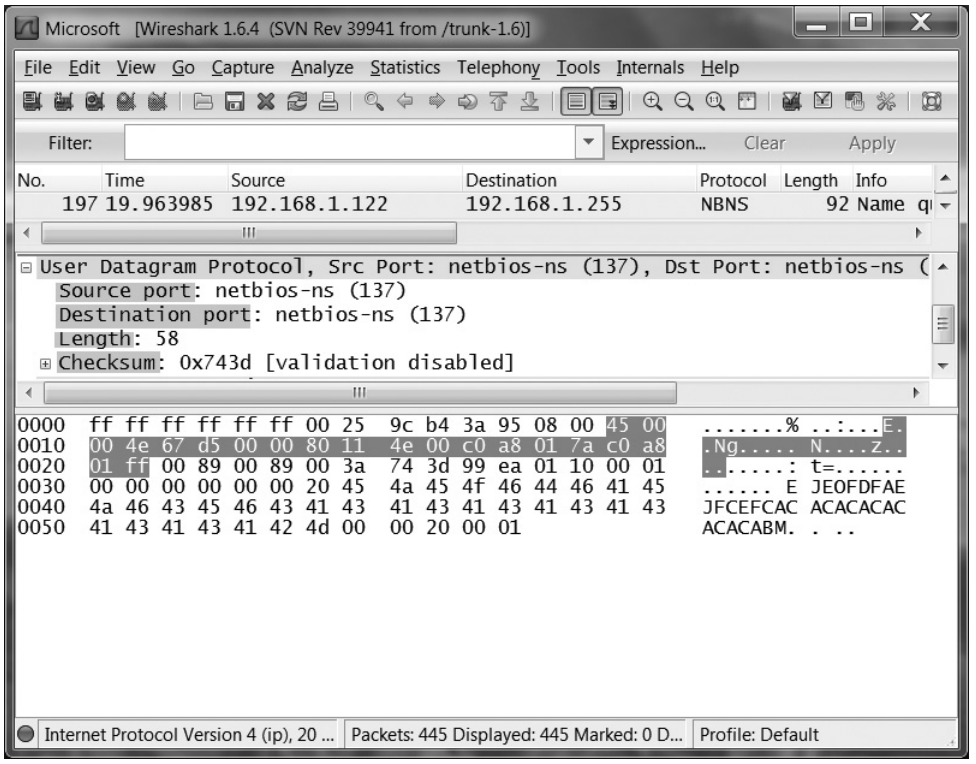

Source: Wireshark

Figure 2-6 UDP header structure as seen in a Wireshark packet capture

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Fragmentation creates a number of security problems, however. Because the TCP or UDP port number is supplied only at the beginning of a packet, it appears only in fragment number 0. Fragments numbered 1 or higher pass through the filter without being scrutinized because they contain no port information. An attacker simply has to modify the IP header to make all fragment numbers start at 1 or higher. All fragments then go through the filter and can access internal resources.

To improve security, you should configure the firewall or packet filter to drop all fragmented packets. You could also have the firewall reassemble fragmented packets and allow only complete packets to pass through. Note that fragmentation is seldom used today because of the improvements in routers.

The TCP Life Cycle and the TCP Three-Way Handshake

Before a client initiates a connection-oriented TCP session with another computer, it must establish which of its own ports it will use as a source of communication and establish the destination port on the other computer. Typically, a client, such as Host A in the following example, will dynamically assign itself a source port on which it will communicate and will know what port is appropriate for the destination port on the other computer. For example, Web servers listen for requests for HTTP service on their port 80; therefore, if Host A were trying to access a Web page on Host B, Host A would set port 80 as the destination port.

To establish connection-oriented communication, each computer needs a way to know that the other computer received the packets sent. Sequence and acknowledgement numbers perform this function, as demonstrated in the way that two hosts first establish the TCP connection: the TCP three-way handshake.

- Host A includes a randomly generated initial sequence number in its first packet to Host B. This packet is called a SYN packet because the TCP SYN flag is set. The acknowledgement number is zero because the SYN packet is the first in the session and there is no previous packet for Host A to acknowledge. Table 2-8 shows parts of the SYN packet configuration for this example.

Sending computer	Host A
Source TCP port	26077
Destination TCP port	80
Sequence number	50088
Acknowledgement number	0
Flags	SYN

© Cengage Learning 2014

Table 2-8 TCP three-way handshake: SYN

- Host B receives the SYN packet and responds with a SYN ACK packet. This packet includes a randomly generated initial sequence number for Host B. As a way of proving that Host B received the SYN packet from Host A, the acknowledgement

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

number is set to the number that Host B expects to receive in the second packet from Host A. The first packet's sequence number is incremented by one and placed as the acknowledgement number. This configuration is shown in Table 2-9.

Sending computer	Host B
Source TCP port	80
Destination TCP port	26077
Sequence number	79995
Acknowledgement number	50089
Flags	SYN ACK

© Cengage Learning 2014

Table 2-9 TCP three-way handshake: SYN ACK

- The final packet in the three-way handshake is the ACK packet that Host A sends in response to the SYN ACK from Host B. Now Host A increments its initial sequence number by one and sets the acknowledgement number to be one more than the initial sequence number that Host B sent in the SYN ACK (see Table 2-10).

Sending computer	Host A
Source TCP port	26077
Destination TCP port	80
Sequence number	50089
Acknowledgement number	79996
Flags	ACK

© Cengage Learning 2014

Table 2-10 TCP three-way handshake: ACK

More goes on in this three-way handshake than acknowledgements. For example, sliding window size is negotiated. After a connection is established, TCP sliding windows control the flow and efficiency of communications. During the transfer of large amounts of data, it would be very inefficient if each packet had to be received and acknowledged before the next packet could be sent. TCP/IP is designed for packet-switching networks where packets might arrive out of order. Sliding window size determines the number of packets that can be sent before ACKs must be received. The sender controls the size of the sliding window.

Once data has been exchanged, either party can end the session by sending a packet with the FIN flag set. The station that receives this initial flag sends a response packet with the ACK flag and its own FIN flag set to acknowledge receipt, and to show that it is also ready to end

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

communications. If the receiving side still has data to send, however, it sends only an ACK flag back and continues sending data until it is done.

After all data has been sent, the side that first received the FIN flag sends its own FIN flag to show that termination can begin. During this lapse between the two flags, the sender of the first FIN flag has a status of FIN WAIT 2, and the recipient of the first FIN flag has a CLOSE WAIT status. Some applications, such as Web browsers, often use this type of half-closed connection to avoid having to initiate the session again. Figure 2-7 summarizes the TCP three-way handshake.

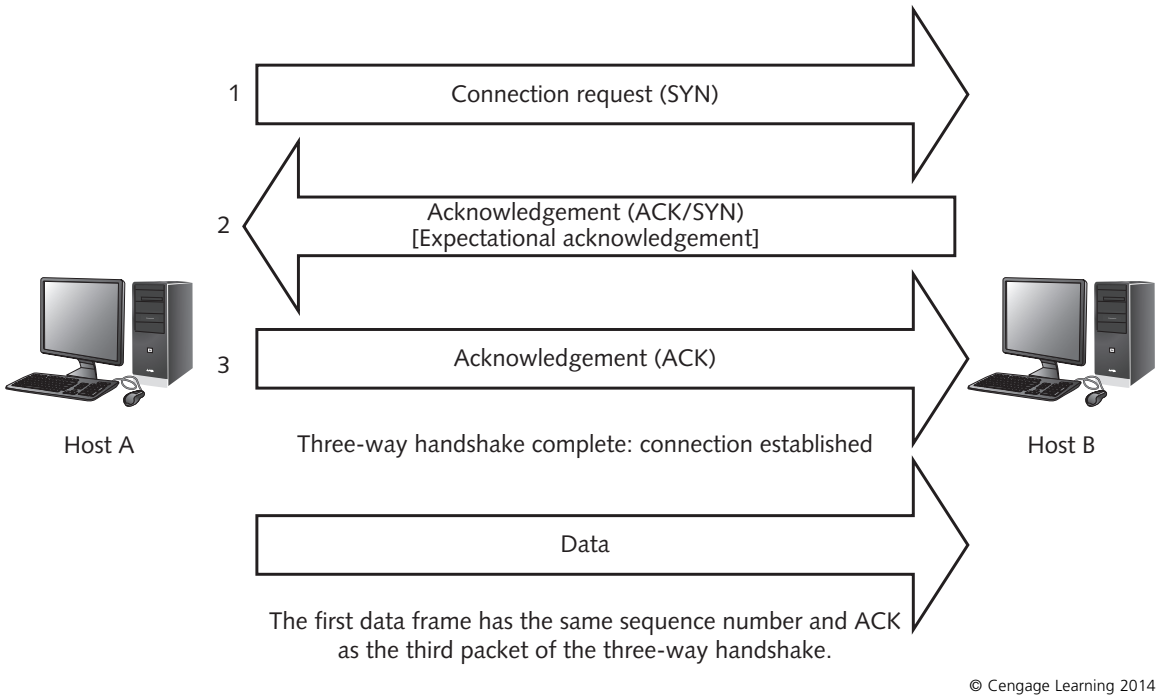

© Cengage Learning 2014

Figure 2-7 Summary of the TCP three-way handshake

Domain Name System

Domain Name System (DNS) is a general-purpose service used both on the Internet and on organizations' internal networks. DNS servers translate **fully qualified domain names (FQDNs)** to IP addresses that identify the host computer. For example, to connect to Web sites, users need access to an authoritative DNS server for the system domain they are trying to reach. If a user enters the FQDN *www.cengage.com*, the authoritative DNS server for cengage.com needs to respond with the correct IP addresses for the computer whose hostname is www.

In terms of network security, DNS is important because it gives network administrators another tool for blocking unwanted communication. With firewalls, Web browsers, and proxy servers, administrators can block DNS names of Web sites and other sites that contain offensive or unsuitable content. Proxy servers are devices that protect internal clients through

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

the use of network address translation. In addition, networks that use DNS servers need to allow DNS traffic when packet filtering is set up.

Attackers can exploit DNS in many ways, including buffer overflow attacks, zone transfer attacks, and cache poisoning attacks. In a DNS buffer overflow attack, an overly long DNS name is sent to the server. When the server is unable to process or interpret the DNS name, it cannot process other requests.

DNS zone files contain a list of every DNS-configured host on a network as well as their IP addresses. Microsoft DNS-enabled networks also list all services running on DNS-configured hosts. When an attacker attempts to penetrate a network, the DNS zone file can provide a list of exploitable targets on that network. When configuring DNS servers connected to the Internet, you should disable zone transfers to all hosts except those that are internal to the network. Internal hosts must be able to transfer zone information to update their records.

A DNS cache poisoning attack exploits the fact that every DNS packet contains a Query section and a Reply section. An older, more vulnerable server has stored answers sent in response to requests to connect to DNS addresses. Attackers can break into the cache to discover the DNS addresses of computers on the network. Most DNS servers, however, have been patched to eliminate this vulnerability.

A newer DNS cache poisoning exploit was discovered by Dan Kaminsky in 2008. This exploit involves the spoofing of transaction IDs, which are supposed to prevent hackers from assigning their own IP addresses to a domain. DNS uses transaction IDs in the range of 0 to 65535. If a hacker sends multiple, slightly varied requests to a name server (such as requests to resolve 1.frog.com, 2.frog.com, and so on), eventually the domain can be spoofed by matching the ID. Once the attacker correctly matches the transaction ID, he can direct all traffic for that site to a site of his choosing. An attacker can also pollute top-level domains using this vulnerability. Prior to releasing details of the exploit, Kaminsky notified vendors and allowed them time to develop a patch that focused on randomizing port numbers.

Originally, the DNS infrastructure did not have optimum security. In 2005, the first implementations of DNSSEC (DNS Security) were rolled out in Sweden in the .se domain. DNSSEC uses cryptographic techniques to enable authentication and data integrity of DNS packets, eliminating vulnerabilities that allow exploitations such as cache poisoning. Unfortunately, DNSSEC is still not widely used, partly because an enormous amount of work is required to revise existing DNS implementations and to establish the complex cryptographic infrastructure. (For more details on cryptography, see Chapter 5.)

Internet Protocol Version 6 (IPv6)

IPv4 has serious drawbacks. Although it was an engineering masterpiece in 1981, the Internet has grown at a rate that the creators of the IPv4 32-bit addressing scheme did not expect. IP addresses are now in short supply, so **Internet Protocol version 6** (**IPv6**), which has a larger address space of 128 bits, is being deployed to allow an almost endless supply of IP addresses.

Because an IPv4 address is 32 bits long, IPv4 permits a total of 2^{32} addresses, which is more than 4 billion. With 128 bits, IPv6 offers 2^{128} addresses, which is 340 undecillion. An undecillion is a 1 followed by 39 zeros.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

IPv4 also presents problems with the routing system. Routers on the Internet backbone have routing tables with about 90,000 entries. Routers get the job done, but because most computers are not connected directly to the Internet backbone, a packet must traverse several extra hops along the route to its destination. In IPv6, backbone routing tables need only the entries of other routers that are connected directly to them. The information in an IPv6 header contains the rest of the information needed to get a packet to its destination, so the process is streamlined.

Security is another concern with IPv4. Although it does support IPsec (an industry standard set of encryption and authentication protocols), IPv4 has no native encryption methods. Plenty of encryption methods are available, but the lack of standardization can create compatibility problems, and encryption can increase overhead on the network. IPv6, on the other hand, has integrated support for IPsec.

Another advantage of IPv6 is that Network Address Translation (NAT) is not needed because of the vast number of IP addresses provided. While NAT has worked well enough to deal with the decreasing number of IP addresses in IPv4, NAT has security problems, as you will see in Chapter 11 when you learn about VPNs. In short, because NAT devices need to read encapsulated IP headers, it is difficult to maintain data confidentiality for end-to-end transmissions; typically, the packets are unencrypted by the NAT firewall and sent through the internal network unencrypted. IPv6 obviates this problem.

Another major advantage of IPv6 is its autoconfiguration capabilities. Instead of relying solely on Dynamic Host Configuration Protocol (DHCP) or manual configuration, IPv6 can determine its own settings based on two different models. **Stateful autoconfiguration**, provided by DHCPv6, is referred to as "stateful" because the DHCPv6 client and server must keep their information updated to prevent addressing conflicts. With **stateless autoconfiguration**, a computer that is trying to connect to a network can determine its own IP address based on its Media Access Control (MAC) address and then can receive additional configurations without a DHCP server. This process simplifies some aspects of network administration because a server does not need to issue address configurations. In the next few sections, you learn about the core IPv6 protocols and see how the next generation of IP works.

IPv6 Core Protocols

You have probably read brief descriptions of IPv6, but most networking books and articles have not explored it in detail until recently. Although you might be a whiz at IPv4, IPv6 has some major differences in its core architecture and functions, and uses some different core protocols. For example, IPv4 uses Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and ICMPv4 Redirect messages to manage node-to-node communications, but IPv6 uses the Neighbor Discovery protocol. The next few sections explain the core architecture and protocols of IPv6.

IPv6 is a connectionless, unreliable datagram protocol used mainly for addressing and routing packets between hosts. Being connectionless, IPv6 does not establish a session before data is exchanged, and delivery is not guaranteed. IPv6 makes a concerted attempt to deliver a packet but relies on higher-layer protocols, such as TCP, if acknowledgement and recovery of lost packets are required.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

An IPv6 datagram consists of the IPv6 header and the IPv6 payload. The IPv6 header is made up of the IPv6 base header and IPv6 optional extension headers. For functional purposes, the optional extension headers and upper-layer protocols are considered part of the IPv6 payload.

As you can see in Figure 2-8, an IPv6 header is more streamlined than an IPv4 header.

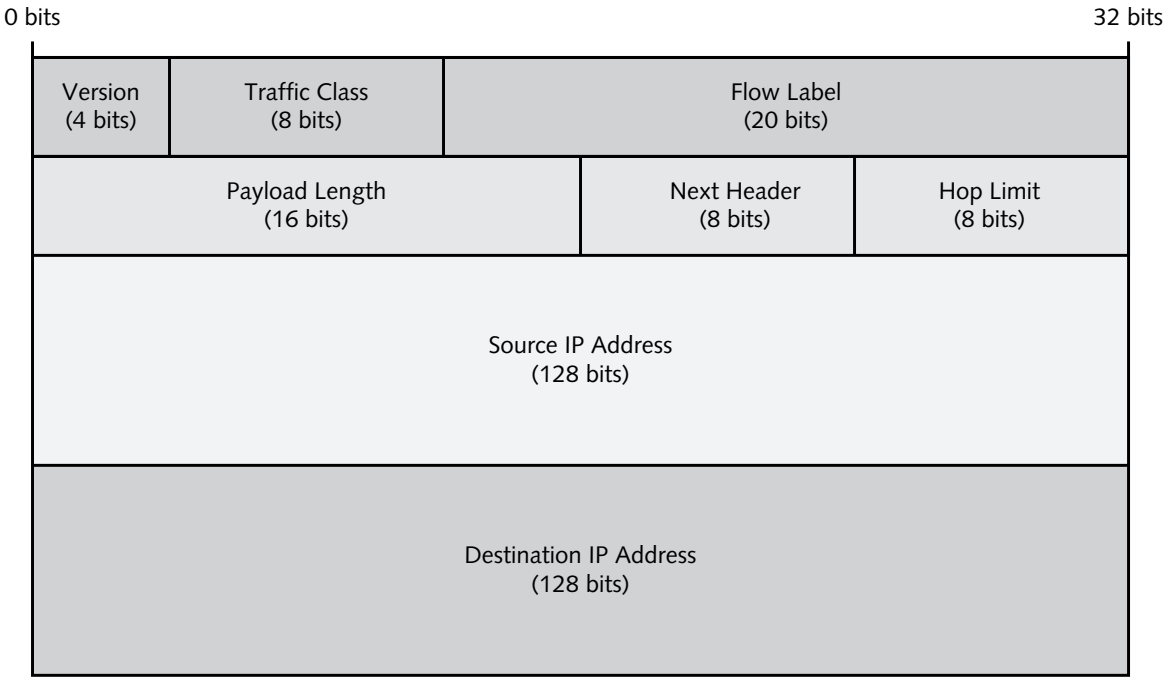

© Cengage Learning 2014

Figure 2-8 IPv6 header structure

The fields in the IPv6 header are as follows:

- *Version*—This field specifies the IP version number (6 for IPv6). Note that the Internet Header Length (IHL) field has been removed in IPv6.
- *Traffic Class*—Also known as the priority field, this field identifies traffic subject to flow control. A value of 1 to 7 indicates lower-priority transmissions that can be slowed down when encountering congestion. Values from 8 to 15 represent real-time traffic that must have a constant sending rate.
- *Flow Label*—This experimental field has not yet been implemented fully, but it can support labeling of some groups of packets as a "flow" and thus improve performance of time-sensitive content such as multimedia. For example, a connection could be specified as "no delay allowed on this connection."
- *Payload Length*—This field indicates the length of the IPv6 payload, which includes the extension headers as well as the upper-layer content. (Extension headers are not shown in Figure 2-8.) The Payload Length field replaces the IPv4 Total Length field

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

and is expressed in 16 bits, indicating a payload of up to 65,535 bytes. If a larger payload is present, this field is set to zero and the Jumbo Payload option is specified in the Hop-by-Hop Options extension header.

- *Next Header*—This field identifies the first extension header. If no extension header is present, the field identifies the name of the transport protocol handler for that packet (for example, TCP or UDP). The port numbers used in the second case are the same as those used in IPv4.
- *Hop Limit*—This field indicates the maximum number of hops the packet is allowed before it is dropped. This number is decreased by 1 each time the packet is transmitted by a router. When the value in this field reaches 0, the packet is dropped. This field replaces the IPv4 Time to Live field.
- *Source IP Address*—This field contains the 128-bit address of the packet source.
- *Destination IP Address*—This field contains the 128-bit address of the packet's intended recipient. It might not be the final destination if a routing extension header is present.

Extension headers are not normally found in a typical IPv6 packet. If needed, however, the sending host adds the appropriate extension header(s). IPv6 extension headers include the following:

- *Hop-by-Hop Options header*—This header carries information that every node along the delivery path must examine and process. It can specify a Jumbo Payload option for payloads greater than 56,535 bytes and up to about 4.2 gigabytes. It also supports the Router Alert option used in Multicast Listener Discovery, which is explained later in this chapter.
- *Destination Options header*—This header carries optional information used by intermediate destinations or by the final destinations. A Next Header value of 60 indicates this header's presence.
- *Routing header*—Similar to IPv4's loose source and record routing (lsrr) option, this header is used to list one or more intermediary nodes (particularly routers) to which the packet is to be forwarded. A Next Header value of 43 indicates that a Routing header is present.
- *Fragment header*—This header is used to send a packet larger than the maximum transmission unit (MTU) value allows. For example, the MTU of Ethernet is 1500 bytes, so larger payloads need to be fragmented. The source node divides large packets into fragments and generates an identification value. (Fragmentation in IPv6 is performed only by source nodes, not by routers along the delivery path.) A Next Header value of 44 indicates a fragment header's presence, but only the destination node processes this header, using it to reassemble fragmented packets on receipt.
- *Authentication header (AH)*—This header provides data authentication and integrity for IPv6 packets. It is algorithm-independent and supports many authentication techniques. The Authentication header is part of the IPsec (IP Security) standard and provides a field for the Security Parameters Index (SPI), which identifies the Security Association (SA) required by IPsec. A value of 51 in the Next Header field specifies the Authentication header.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *Encapsulating Security Payload (ESP) header*—This field provides data integrity, authentication, and confidentiality to the encapsulated payload (the data portion of the datagram). Also, a foundational protocol of IPsec, the ESP header, includes an SPI field.

Remember that IPv4 and IPv6 headers are not interoperable. A host or router must be configured to support both so that it can recognize and process both formats.

Internet Control Message Protocol for IPv6 ICMPv6, an integral component of IPv6 communications, is used by IPv6 nodes for reporting errors and for diagnostic purposes. As in ICMPv4, ICMPv6 uses the ping and tracert commands as well as other diagnostics you already know.

ICMPv6 messages are grouped into two classes: error messages and informational messages. Error messages are identified by message type codes 0 to 127, and informational messages are identified by message type codes 128 to 255. Table 2-11 shows common message type codes for ICMPv6.

Message type	Type
Destination Unreachable	1
Packet Too Big	2
Time Exceeded	3
Parameter Problems	4
Echo Request	128
Echo Reply	129

© Cengage Learning 2014

Table 2-11 Common ICMPv6 message type codes

ICMPv6 is specified in RFC 4443 and updated by RFC 4884.

An ICMPv6 message is preceded by an IPv6 header and sometimes by extension headers. An ICMPv6 header follows the IPv6 header and is identified by a Next Header value of 58 in the IPv6 header. An ICMPv6 header has the format shown in Figure 2-9.

The Type field in Figure 2-9 contains the value for a type of ICMPv6 message; for example, a value of 128 indicates an Echo Request message. The Code field specifies additional parameters for the message. For example, if a Destination Unreachable (type 1) message is received, the Code field might contain 1, indicating that the firewall configuration or other security protocol denied packet access, thus prohibiting communication with the destination. The Checksum field is used to detect data corruption in the ICMPv6 message and parts of the IPv6 header.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

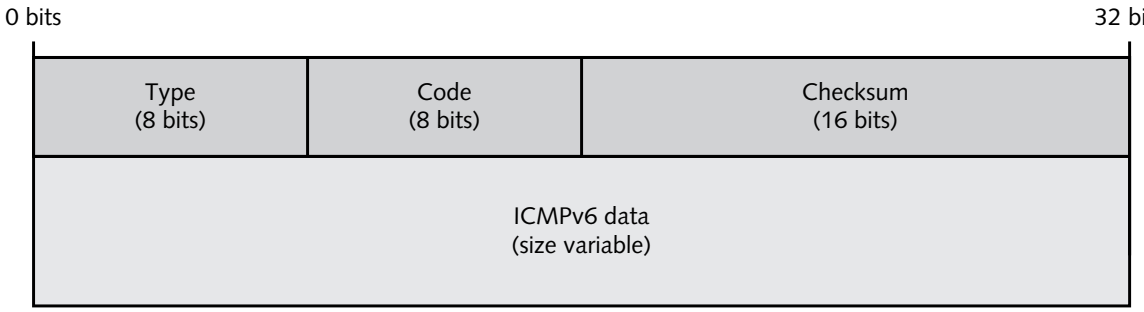

© Cengage Learning 2014

Figure 2-9 ICMPv6 header structure

ICMPv6 also provides the framework for additional IPv6 features, as shown in Table 2-12. The following sections discuss these features in more detail.

Message type	Code
Multicast Listener Discovery (MLD)	MLD replaces IGMPv3 (Internet Group Message Protocol), which is used for controlling multicasts in IPv4. MLD is a series of three ICMPv6 messages used to manage subnet multicast membership. (See Table 2-13.)
Neighbor Discovery (ND)	ND replaces ARP, ICMPv4 Router Discovery, and ICMPv4 Redirect and has other functions, including prefix detection, duplicate address detection, and automatic address configuration. ND is a series of five ICMPv6 messages used to manage node-to-node communications on a link. (See Table 2-14.)

© Cengage Learning 2014

Table 2-12 ICMPv6 features

Multicast Listener Discovery Multicasts are used for a variety of network functions and connectionless delivery of information to multiple subscribers at the same time. Unlike conventional data streaming, which uses one stream per recipient, multicasting has a single stream on any link with at least one recipient. Instead of consuming bandwidth to distribute multiple copies, routers track subscribers (group members) and create copies only as needed.

IP multicast traffic is sent to a single address but is processed by all members of a multicast group. Hosts listening on a specific multicast address are part of the multicast group. Group membership is dynamic, with hosts joining and leaving the group at any time. A host does not need to be a group member to send messages to the group. The group size is unlimited, and members can be on different subnets as long as the connecting routers support multicast message forwarding between those subnets. Multicast group members can be members of multiple groups simultaneously.

Multicast addresses can be permanent or transient. A permanent multicast address has an administratively assigned IP address. However, the address is permanent, not the group members. In IPv6, **Multicast Listener Discovery** (**MLD**) enables IPv6 routers to discover multicast listeners on a directly connected link and decide which multicast addresses are of interest to the nodes. MLD uses a series of ICMPv6 messages to track membership, as shown in Table 2-13.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

MLD message type	Description
Multicast Listener Query	Multicast routers send queries to poll a network segment for group members. Queries can be general, can request membership for all groups, or can be used for a specific group.
Multicast Listener Report	This message is sent by a host when it joins a multicast group or in response to a Multicast Listener Query.
Multicast Listener Done	This message is sent by a host when it leaves a host group and is the last member of that group on the network segment.

© Cengage Learning 2014

Table 2-13 Multicast Listener Discovery message types

MLD for IPv6 is specified in RFC 2710, with updates added in RFCs 3590 and 3810. Be sure to check for other updates because Internet drafts change often.

Neighbor Discovery In IPv4, ARP is used to resolve IP addresses to MAC (Media Access Control) addresses, and ICMP Router Discovery and ICMP Redirect are used to locate neighboring routers and redirect hosts to better routes to reach destination addresses. IPv6 uses a new protocol, **Neighbor Discovery (ND)**, to handle these tasks and provide additional functions. ND uses ICMPv6 messages to manage node-to-node communications. Table 2-14 summarizes the functions of ND.

Process	Description
Router discovery	Discovers neighboring routers
Prefix discovery	Discovers local network prefixes (equivalent to IPv4 network addresses)
Parameter discovery	Discovers additional parameters, such as MTU (Maximum Transmission Unit) size for the network segment and default hop limit for outbound packets
Address autoconfiguration	Automatically configures addresses
Address resolution	Resolves a neighboring node's address to its MAC address
Next-hop determination	Determines the next-hop node address, which is typically the final destination or a router on the network segment
Neighbor unreachability detection	Determines whether neighboring hosts or routers are no longer available
Duplicate address detection	Determines that an address considered for use is not already in use by a neighboring node
Redirect function	Determines the process by which a router informs a host of a better first-hop IPv6 address to reach a destination

© Cengage Learning 2014

Table 2-14 IPv6 Neighbor Discovery functions

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

ND for IPv6 is specified in RFC 4861, with updates added in RFC 4311 and RFC 5942.

ND defines five different types of ICMP messages:

- Router Solicitation messages are sent by hosts when an interface is enabled so that they can discover routers on the network. A Router Solicitation message requests that routers send a Router Advertisement message immediately rather than at the next scheduled time.
- Router Advertisement messages inform hosts about router presence and provide additional parameters about the link or services, such as address configuration or suggested hop limits. A Router Advertisement message is sent at defined intervals or in response to a Router Solicitation message.
- Neighbor Solicitation messages are sent by a node to determine the MAC address of a neighbor or to verify that a neighbor is still reachable. Neighbor Solicitation messages are also used for duplicate address detection.
- Neighbor Advertisements are sent in response to a Neighbor Solicitation message or to update neighbors of a MAC address change.
- Redirect messages are sent by routers to tell hosts about better first-hop addresses to reach a destination.

ICMP is an efficient method of managing many underlying networking functions, such as those formerly assumed by Internet Group Management Protocol (IGMP) and ARP. ND and MLD take care of these housekeeping jobs, such as managing group membership and managing node-to-node communications. Each message is identified in the packet header by the corresponding ICMPv6 message type code, as shown in Table 2-15. (In addition to these ICMPv6 message types, remember the common message types in Table 2-11.)

Message type	Type
Group Membership Query (MLD)	130
Group Membership Report (MLD)	131
Group Membership Reduction/Done (MLD)	132
Router Solicitation (ND)	133
Router Advertisement (ND)	134
Neighbor Solicitation (ND)	135
Neighbor Advertisement (ND)	136
Redirect (ND)	137

© Cengage Learning 2014

Table 2-15 Multicast Listener Discovery and Neighbor Discovery message types

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

IPv6 Addressing

An IPv6 address is 128 bits long, which would mean a very long number if you used the same decimal numbering scheme as in IPv4—not to mention using binary and writing out 128 ones and zeros. To make IPv6 addresses manageable, the hexadecimal numbering format known as base 16 is used. (The hexadecimal format is often called *hex*.) An IPv6 address consists of eight hex groups separated by colons. Each hex group contains a 16-bit value, and each digit represents a 4-bit value. The following examples show what an IPv6 address looks like:

4EDC:0000:7654:3210:F3DC:BA98:7654:AB1F

1080:0:0:0:8:800:200C:417A

Including leading zeros in a group is not necessary, and hex letters are not case sensitive. You can also replace consecutive zeros with a double colon, as shown in this modification:

1080::8:800:200C:417A

This "compression" of leading zeros comes in handy when typing long strings, but remember that the double colon can be used only once in an address. In a mixed environment of IPv6 and IPv4, you can use the colon hexadecimal notation of IPv6 addresses and the 32-bit dotted decimal notation of IPv4 addresses, as shown in the following examples:

0:0:0:0:0:0:22.1.68.45

0:0:0:0:0:FFFF:131.123.2.8

Using the double colon, you can condense these addresses to the following:

::22.1.68.45

::FFFF:131.123.2.8

You might wonder what address is used for the loopback. It looks like 0:0:0:0:0:0:0:1, which can be compressed to ::1.

Unicast, Multicast, and Anycast Addressing

IPv6 uses three types of addresses: unicast, multicast, and anycast. Notice that IPv6 does not use broadcast addresses.

Unicast is used for one-to-one communication, such as that between two single hosts. Another example is communication between two routers. A unicast address is configured for each interface connected to the network. IPv6 has several forms, or **scopes**, of unicast addresses:

- *Global unicast addresses*—Equivalent to public addresses in IPv4, these addresses are routable on the Internet.
- *Site-local unicast addresses*—Similar to private addresses in IPv4, this type of address is being phased out and is not permitted in new IPv6 implementations.
- *Unique local IPv6 unicast addresses*—These addresses are replacing the sometimes ambiguous site-local unicast addresses. Unique local IPv6 unicast addresses are private to an organization but are still unique throughout the organization.
- *Link-local unicast addresses*—These addresses are used by hosts when they communicate with other hosts on their same network segment. These addresses are equivalent to the APIPA (Automatic Private IP Addressing) used in IPv4. Link-local unicast addresses always begin with FE80 in the first 16-bit section.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Multicast is used for one-to-many communication, in which a single host can send packets to a group of recipients. Multicast addresses also use scopes: site-local, link-local, and interface-local. Multicast addresses always begin with FF in the first byte.

Anycast addresses are used for one-to-one or one-to-many communication. Anycast addresses are not assigned a specific range; instead, they are created automatically when a unicast address is assigned to more than one interface. This means that anycast addresses are assigned from unicast address ranges and have the same scopes as unicast addresses. The idea behind anycast is to offer flexibility in providing services. For example, a group of servers might provide a service. Using anycast, the service can be provided by any one of the servers—usually the one that is closest. Anycast addresses are currently used only by routers, but their use will expand as the technology becomes widespread.

IPv6 Configuration

Microsoft operating systems since Windows XP SP 1 have built-in IPv6 support. These platforms support stateless autoconfiguration and do not usually need manual configuration. By default, a link-local address is assigned to every Ethernet interface during startup. IPv6 addresses, such as global addresses, are assigned automatically based on the receipt of IPv6 Router Advertisement messages. You must have a correctly configured IPv6-capable router on your network segment to receive additional addresses through IPv6 Router Advertisement messages. Note that a single host often has more than one IPv6 address. For example, a host with a global address will also have a link-local address.

Some manual configuration is required for more advanced features and setups of IPv6. You can find procedures for your equipment with an Internet search or by contacting the vendor for instructions.

IPv6 Utilities

IPv6 includes several integrated utilities for configuration, troubleshooting, and other administrative tasks. Some utilities are familiar because they are used in IPv4 and have not changed much. In the following sections, you examine some major Windows IPv6 utilities.

IPv6 supports many different tools, and an Internet search for IPv6 tools yields many sites that offer handy utilities, including *www.ipv6tools.org*, *www.ultratools.com/ipv6Tools*, and *http://dns.antd.nist.gov/ipv6/*.

Ipconfig The ipconfig command shows IPv6 configuration details when IPv6 is installed on a Windows operating system. You can also use the command with IPv4. Figure 2-10 shows the result of using the ipconfig command on a system running a dual stack (both IPv4 and IPv6). Note that the host has both a link-local address and a global IPv6 address on the tunnel adapter and a different link-local address on the Ethernet interface, demonstrating that hosts in IPv6 can have multiple IPv6 addresses.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

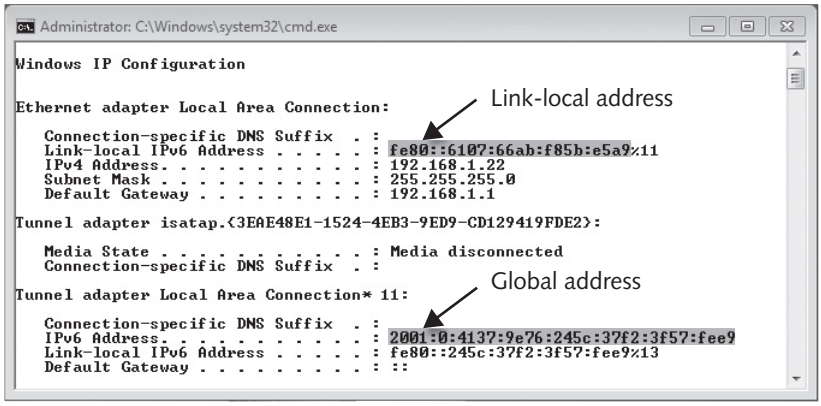

Source: Used with permission of Microsoft Corporation

Figure 2-10 Using the ipconfig command

Netstat You can display the system's routing tables by using the netstat -r command, as shown in Figure 2-11. Notice that this example includes routes for three scopes: global, link-local, and multicast.

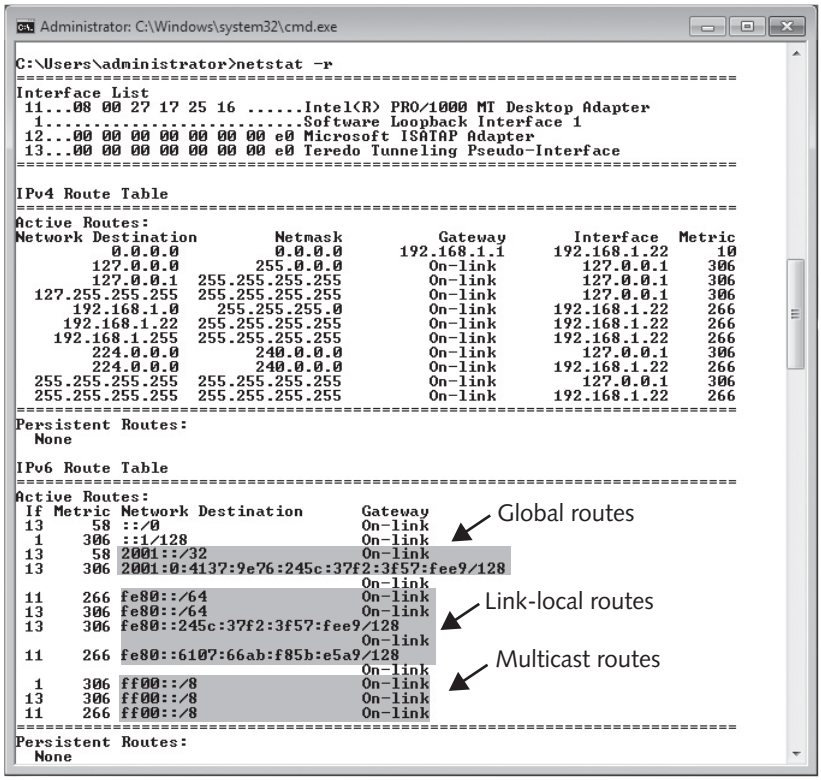

Source: Used with permission of Microsoft Corporation

Figure 2-11 Using the netstat –r command

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Use the netstat command with the -n option to show current sessions with the associated port numbers. Figure 2-12 shows an IPv6 session between port 49232 on the Windows 7 local machine and its Windows Server 2008 domain controller's port 445.

Source: Used with permission of Microsoft Corporation

Figure 2-12 Using the netstat command

Use the command netstat -ps IPv6 to display detailed statistics on IPv6 activity since the last bootup. To get even more detailed information, use commands like netstat -ps TCPv6, netstat -ps UDPv6, and netstat -ps ICMPv6.

Netsh Netsh is a command-line scripting tool on Windows systems that allows troubleshooting and configuration of network interfaces. Netsh works in both noninteractive and interactive modes. In the first of these modes, you can enter commands at the command prompt and receive results. In interactive mode, you can enter the Netsh utility by entering the netsh command, and then navigate the different contexts of the netsh tools.

Chapter Summary

- Transmission Control Protocol/Internet Protocol (TCP/IP) is a suite of many protocols for transmitting information from point to point on a network.
- The subprotocols of TCP/IP can be mapped to the layers of the Open Systems Interconnection model (OSI). TCP and UDP map to the Transport layer. IPv4, IPv6, ICMP, and ICMPv6 map to the Network layer.
- The IP addresses most commonly used on the Internet conform to IPv4, with addresses of 32 bits (4 bytes) of data. An IP address consists of two main parts: a network identifier and a host identifier. The two values are distinguished by a subnet mask.
- Classless Interdomain Routing (CIDR) is an address notation scheme that specifies the number of masked bits in an IP address/subnet mask combination.
- TCP/IP is transmitted along networks as discrete chunks called segments, datagrams, packets, or frames. These terms are not interchangeable. Each IP datagram is divided into header and data sections. The data in an IP packet is the part that users see along with upper-layer protocol headers, but the IP header is the part that computers use to communicate.
- You must understand the normal configuration of the fields in IP, TCP, and UDP headers to recognize and filter unwanted or malicious traffic.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- IP networks can be subnetted into smaller network segments that permit more efficient use of IP addresses, reduction of network traffic, and increased network security.
- Fragmentation of IP packets was originally developed so that large packets could pass through routers with frame size limitations.
- Domain Name System (DNS) translates fully qualified domain names, such as *www.cengage.com*, into IP addresses.
- Before a client initiates a TCP session, it must determine the port number that identifies the session and the starting sequence number.
- The TCP three-way handshake establishes a reliable connection between two points.
- IPv4 presents problems with the current routing system, requiring routers on the Internet backbone to store tens of thousands of entries. Security is another concern with IPv4. Although it does support IPsec, IPv4 has no native encryption methods.
- IPv6 was designed to address some problems with IPv4. IPv6 has a much larger address space of 128 bits and native support for IPsec, and the addressing architecture requires routers to store routing table entries only for nodes connected directly to them. The IPv6 packet contains all other information required for delivery.
- IPv6 is a connectionless, unreliable datagram protocol used mainly for addressing and routing packets between hosts. An IPv6 datagram consists of the IPv6 header and IPv6 payload. The IPv6 header has been streamlined considerably more than IPv4 headers. IPv4 and IPv6 headers are not interoperable.
- ICMP is used for reporting errors and for diagnostic purposes. As in ICMPv4, ICMPv6 uses the ping and tracert commands as well as other familiar diagnostics. ICMPv6 uses error messages and informational messages. ICMPv6 also enables two new protocols in IPv6: Multicast Listener Discovery (MLD) and Neighbor Discovery (ND).
- MLD enables IPv6 routers to discover multicast listeners on a directly connected link and to decide which multicast addresses are of interest to those nodes.
- IPv6 uses ND to perform the tasks that ARP, ICMP Router Discovery, and ICMP Redirect handled in IPv4. ND uses ICMPv6 messages to manage node-to-node communications.
- The hexadecimal numbering format makes IPv6 addresses manageable. An IPv6 address consists of eight hex groups separated by colons. Each hex group contains a 16-bit value, and each digit represents a 4-bit value.
- IPv6 uses three types of addresses: unicast, multicast, and anycast. IPv6 does not use broadcast addresses.
- IPv6 supports stateless address autoconfiguration. By default, a link-local address is assigned to every Ethernet interface during startup. Site-local addresses or global addresses are assigned automatically based on receipt of IPv6 Router Advertisement messages.
- You can monitor and configure IPv6 using tools such as Ipconfig, Netstat, and Netsh.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Key Terms

anycast An address created automatically when a unicast address is assigned to more than one interface. Anycast addresses are assigned from unicast address ranges and have the same scopes as unicast addresses.

broadcast A communication sent to all hosts on a specific network.

Classless Interdomain Routing (CIDR) An IP address notation method that uses a slash (/) followed by the number of masked bits for an address—for example, 192.168.6.5/27 instead of an IP address of 192.168.6.5 and a subnet mask of 255.255.255.224.

datagram A discrete chunk of information; each datagram contains source and destination addresses, control settings, and data.

Domain Name System (DNS) A network service that provides a database for the resolution of fully qualified domain names to IP addresses, and of IP addresses to fully qualified domain names.

footer A section sometimes added to a TCP/IP packet to indicate the end of the packet.

fragmentation The division of packets into smaller sizes to accommodate routers with frame size limitations.

fully qualified domain names (FQDNs) Complete DNS names of computers that include the computer name, domain name, and domain name extension, such as *www.cengage.com*.

header The part of a packet that contains source and destination information and general information about the packet.

host identifier The part of an IP address that is unique to a computer in its subnet.

Internet Control Message Protocol (ICMP) A protocol that reports network communication errors to support IP communications. The ping command is a common troubleshooting utility based on ICMP.

Internet Protocol version 4 (IPv4) The IP addressing system currently in widespread use on the Internet, in which addresses are created with 32 bits (4 bytes) of data.

Internet Protocol version 6 (IPv6) A new version of IP that is gaining support among software and hardware manufacturers, and that will eventually replace IPv4; this version calls for 128-bit IP addresses.

multicast A transmission used for one-to-many communication, in which a single host can send packets to a group of recipients.

Multicast Listener Discovery (MLD) A core IPv6 protocol that enables IPv6 routers to discover multicast listeners on a directly connected link and to decide which multicast addresses are of interest to those nodes.

Neighbor Discovery (ND) A core IPv6 protocol used to resolve addresses, locate neighboring routers, and redirect hosts to better routes to reach destination addresses. ND uses ICMPv6 messages to manage node-to-node communications.

Network Address Translation (NAT) A process by which internal hosts are assigned private IP addresses and communicate with the Internet using the public address of the NAT device.

network identifier The part of an IP address that a computer has in common with other computers in its subnet.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

scopes Unicast addresses used in IPv6 to identify the application suitable for the address; scopes include global unicast, site-local unicast, and link-local unicast.

stateful autoconfiguration In IPv6, the method of dynamic host configuration that uses DHCPv6. The method is called "stateful" because the DHCPv6 client and server must keep their information updated to prevent addressing conflicts.

stateless autoconfiguration A feature of IPv6 in which a computer can connect to a network by determining its own IP address based on the addressing of neighboring nodes.

subnet mask A value that defines which part of a computer's IP address is its network identifier and which part is the host identifier.

Transmission Control Protocol/Internet Protocol (TCP/IP) A suite of protocols for transmitting information from point to point on a network.

unicast A transmission in which one packet is sent from a server to each client that requests a file or application.

User Datagram Protocol (UDP) A core transport protocol of the TCP/IP suite. UDP is connectionless, meaning it does not ensure delivery or provide ordering, as TCP does. UDP is much faster, and is useful for transmissions that require speed over reliability. UDP relies on upper-level protocols for error-checking and sequencing services.

variable length subnet masking (VLSM) A means of allocating IP addressing according to the network's needs that involves applying masks of varying sizes to the same network. This method creates subnets within subnets and multiple divisions of an IP network.

Review Questions

1. What advantages does IPv6 have over IPv4? (Choose all that apply.)
   a. IPv6 uses DHCP for its configuration settings.
   b. IPv6 uses a 128-bit address space.
   c. IPv4 cannot support IPsec.
   d. IPv6 incorporates IPsec.

2. A Class C address has a first octet decimal range of _____ to _____.
   a. 172, 191
   b. 191, 224
   c. 192, 239
   d. 192, 223

3. Which of the following is a method of hiding internal host IP addresses? (Choose all that apply.)
   a. Network Address Translation (NAT)
   b. configuring the computer to insert a fake source IP address into outgoing messages
   c. proxy servers
   d. setting up software firewalls on all internal hosts

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

4. The Class A address 127.0.0.1 is used for which of the following?
   a. broadcasting to all hosts on a subnet
   b. testing the local TCP/IP software implementation
   c. experimentation
   d. testing the local NIC

5. Why is UDP considered unreliable?
   a. The header does not contain a checksum.
   b. The data is transmitted in clear text.
   c. It is connectionless.
   d. Routers typically drop a large number of UDP packets.

6. In CIDR notation, the IP address and subnet mask 191.9.205.22 255.255.192.0 are written as _____.
   a. 191.9.205.22/19
   b. 191.9.205.22/18
   c. 191.9.205.22/17
   d. 191.9.205.22/16

7. How do routers handle packets that are too large to pass through because of frame size limitations?
   a. Routers drop packets that are too large.
   b. Routers bounce packets back to the sender to be resized.
   c. Routers adjust their MTUs to accommodate the oversized packet.
   d. Routers break packets into smaller pieces called fragments.

8. Which of the following is an IPv6 protocol? (Choose all that apply.)
   a. Multicast Listening Detection
   b. IGMPv6
   c. Multicast Listener Discovery
   d. Neighbor Discovery

9. A DNS server translates _____ to _____.
   a. encrypted IP addresses, clear text
   b. IP addresses, MAC addresses
   c. FQDNs, IP addresses
   d. static addresses, DHCP

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

10. Why is fragmentation considered a security risk?
    a. Fragments numbered 0 contain port information.
    b. Fragments numbered 1 or higher are passed through filters.
    c. Fragmented packets cannot be assembled.
    d. Fragmentation is frequently used.

11. Which of the following is used for one-to-many communication, in which a single host can send packets to a group of recipients?
    a. multicast
    b. unicast
    c. anycast
    d. netcast

12. The number of TCP segments that can be sent before an acknowledgement must be received is determined by the _____.
    a. sequence number
    b. sliding window size
    c. transmission rate
    d. port number in use for the session

13. When one host wants to initiate a TCP session with another host, it sends a packet with the _____ flag set.
    a. SYN
    b. ACK
    c. RST
    d. FIN

14. An ICMPv6 header is indicated by a Next Header value of _____.
    a. 60
    b. 54
    c. 58
    d. 22

15. Compressing the IPv6 address 1080:0:0:0:8:800:200C:417A results in which of the following?
    a. 1080::8:8:::2::C:417A
    b. 1080::8:800:200C:417A
    c. 1080::8:8::::20:C:417A
    d. :1080::8:800:200C:417A

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Projects

Hands-On Project 2-1: Installing the Wireshark Protocol Analyzer

Time Required: 10 minutes

Objective: Download and install Wireshark.

Description: You access the Wireshark Web site, and then download and install Wireshark on both the Windows 7 and Windows Server 2008 computers.

1. Log on to Windows 7 with an administrative account. Log on to Windows Server 2008 with an administrative account. On both systems, complete the following steps.
2. Start your Web browser, and go to **www.wireshark.org/download.html**.
3. Click the appropriate **Windows Installer** link, depending on whether you are using a 32-bit or 64-bit operating system. If the Internet Explorer warning appears about trusted sites, click **Add**, click **Add** again, click **Close**, and then click the **Windows Installer** link again.
4. In the download window, click **Run**. In the next window, click **Run** again. In the Welcome window, click **Next**. In the License Agreement window, click **I Agree**. In the Choose Components window, click **Next**. In the Select Additional Tasks window, click the box next to **Desktop Icon**, and click **Next**. In the Choose Install Location window, click **Next**.
5. Wireshark requires WinPcap. If you completed the Hands-On Projects in Chapter 1, WinPcap is already installed on Windows Server 2008, so you can click **Install** and skip to Step 6. If you did not complete the Chapter 1 Hands-On Projects on Windows Server 2008, or if you are on the Windows 7 system, complete the rest of this step. Verify that the box next to **Install WinPcap** is checked, and click **Install**. In the WinPcap window, click **Next**. In the Welcome to the WinPcap Setup Wizard, click **Next**. Click **I Agree** and click **Install**. Click **Finish**.
6. In the Installation Complete window, click **Next**. In the Completing the Wireshark Setup Wizard, click **Finish**.
7. Leave the systems logged on for the next Hands-On Project.

Hands-On Project 2-2: Using Wireshark to Capture IPv4 Pings

Time Required: 20 minutes

Objective: Capture and begin to analyze network traffic.

Description: To get a better idea of what TCP/IP packet headers look like, you can use a network traffic analyzer to capture packets as they enter or leave your network. In this activity, you capture IPv4 ping packets with Wireshark and begin to analyze how ICMP and ARP traffic function.

1. Log on to Windows Server 2008 and Windows 7 with an administrative account.
2. On both systems, access a command prompt and enter **ipconfig** to verify the systems' IP addresses. Verify that you have connectivity by pinging the IPv4 address of Windows

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Server 2008 from Windows 7. If this ping is not successful, you need to troubleshoot connectivity before continuing with the project.

3. On both systems, open Wireshark. Click **Interface List** in the upper-left corner to display the interfaces detected by Wireshark. Figure 2-13 shows the Capture Interfaces window, where only one interface has been detected. If multiple interfaces are displayed, use the one that shows the most activity in the Packets column.

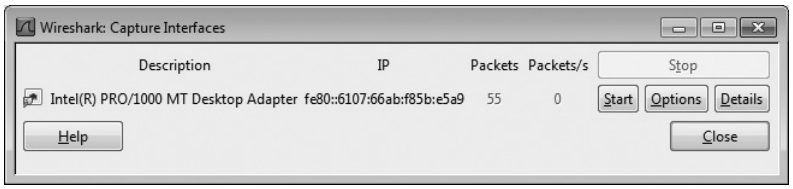

Source: Wireshark

Figure 2-13 Wireshark Capture Interfaces window

4. Before you start to capture network traffic, open a command prompt and enter a ping to the IPv4 address of your partner's system, but do not press Enter yet. Arrange your desktop so that you can see both the command prompt window and the Wireshark Capture Interfaces window. You should perform the next few steps quickly so that you do not capture excess traffic. Click the **Start** button next to the interface with the most activity. The Capture Interfaces window closes and the Wireshark window opens. Immediately click in the command prompt window and press **Enter** to begin the ping. When the ping is complete, immediately click the **Capture** menu in Wireshark and click **Stop**. Each partner should perform these steps.
5. Figure 2-14 shows a typical capture of the preceding steps. Compare your capture to Figure 2-14 and note the following explanations. Notice the entries in the Protocol column; the ARP and ICMP protocols are being used. Each line in the upper frame of the window represents one packet. ARP is used to resolve IP addresses to MAC addresses. Notice the Source and Destination columns. In Figure 2-14, the first frame came from a MAC address that ended in 17:25:16 and was sent as a broadcast to every host on the network. In the Info column, you can see that the sender of the packet wanted to know who has an IP address of 192.168.1.110. The sender wants that system to tell 192.168.1.22 its MAC address. In the second packet, you can see that the source and destination addresses are also MAC addresses, but that no broadcast was needed. In the Info column, you can see that the server answered and provided its MAC address.
6. In the third frame, the protocol changes to ICMP, the protocol that ping uses. The Info section indicates that the frame is an Echo request, the first step of a ping. The Source column indicates that the frame came from 192.168.1.22. But why do more ARPs follow the third packet? Why doesn't the Windows Server 2008 computer simply send an Echo Reply frame? In fact, the Windows Server 2008 computer does not know the MAC address of the Windows 7 computer, and the ARP processes must be reversed so that the Windows Server 2008 computer can learn which MAC address to use to send the Echo Reply.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

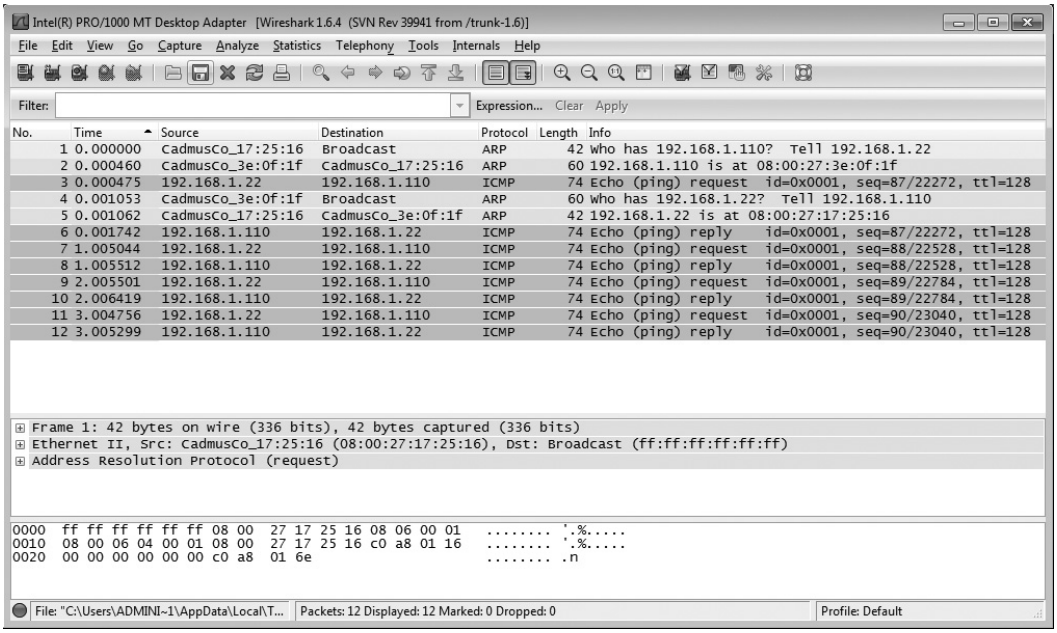

Source: Wireshark

Figure 2-14 Wireshark capture

7. In subsequent packets, you can see that the four Echo Requests and Echo Replies are completed.
8. Leave your Wireshark ARP and ICMP capture open, and leave the systems logged on for the next project.

Hands-On Project 2-3: Examining Individual ARP and Ping Packets

Time Required: 25 minutes

Objective: Learn how Wireshark presents captured packets.

Description: In the previous project, you captured ARP and ICMP packets and examined the functional behavior of the packets from a high-level view. In this project, you open the individual packets and learn how Wireshark presents headers and the data contained in them. You can complete this project on both the Windows 7 and Windows Server 2008 computers.

1. Your capture from Hands-On Project 2-2 is probably not as clean as that shown in Figure 2-14. You may have picked up packets from another activity on your network. In this project, you need to identify the equivalent packets in your capture that match those discussed in these steps. To make this process easier, record the configuration of your Windows 7 and Windows Server 2008 systems by filling in Table 2-16. You can get this information by running the ipconfig /all command on each system.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

	Windows 7	Windows Server 2008
IPv4 address		
IPv6 address		
MAC address		

© Cengage Learning 2014

Table 2-16 Windows systems configuration

The following explanations refer to Figure 2-14 and are described from the perspective of Windows 7 pinging Windows Server 2008, but the student working on Windows Server 2008 can simply reverse the addresses. Of course, your addresses will be different from those in the author's lab setting.

2. Click the first ARP frame related to this activity in the upper part of the Wireshark window. In Figure 2-15, this is the first packet, as indicated in the top row of the No. column. The selected frame changes to blue and the data changes in the middle and bottom frames of the Wireshark window. Notice that you can resize each of the three Wireshark window frames horizontally. Figure 2-15 shows the first packet highlighted.

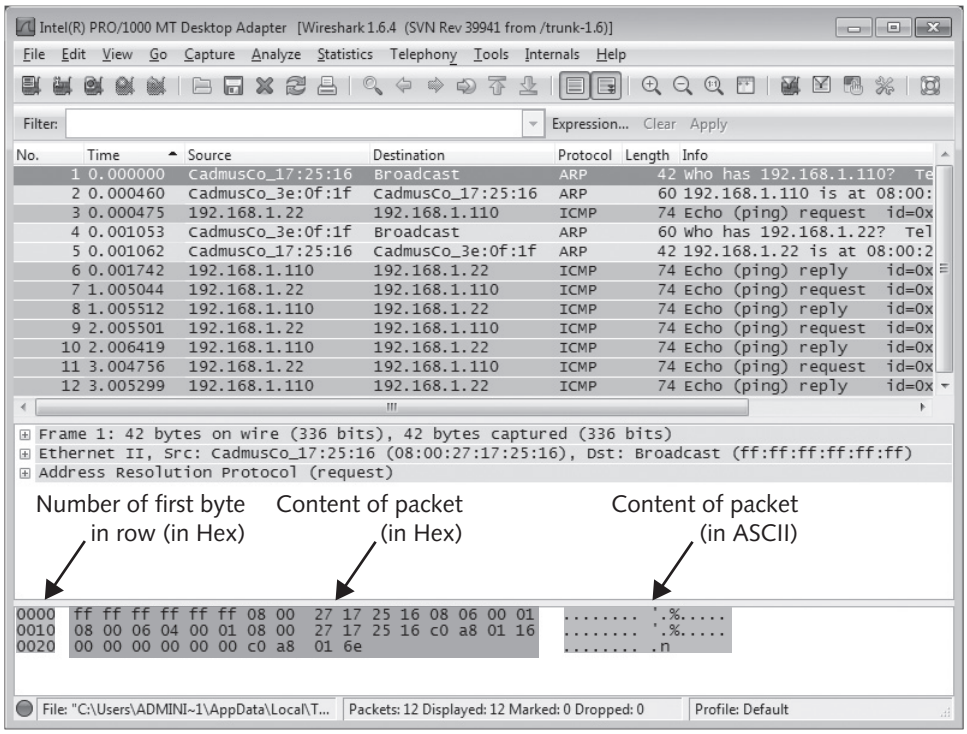

Source: Wireshark

Figure 2-15 First ARP packet selected

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

3. As you have seen, the top frame displays packets in the order they were captured. The Time column tracks elapsed time from the beginning of the capture session. You have already seen how the Source and Destination columns can list MAC or IP addresses depending on the function of the packet. You have also examined the Protocol and Info columns. The only other column is Length, which shows the length of the frame in bytes. Look at the bottom frame and notice the three main columns labeled in Figure 2-15. The first column indicates which byte is the first one shown in the middle Hex column. The first byte in the packet is "ff," as shown in the first byte of the Hex column. (In fact, the first 6 bytes are "ff.") Therefore, the "0000" in the first column indicates that the first byte in the first Hex row is byte zero. The first byte in the second Hex row, "08," is the 16th byte in the packet. The right column in the bottom frame displays the content of the packet in ASCII (American Standard Code for Information Interchange) notation. The packet in Figure 2-15 does not display much useful ASCII information, but you will see some interesting content in the ASCII column later.
4. Click the **+** button to the left of the Frame row in the middle frame of the Wireshark window. The section expands to display information that Wireshark collected about the packet. You may have to resize the middle frame to see all of the content. (Note that Wireshark considers the packets to be frames.) Also, notice the line that reads [Coloring Rule Name: ARP]. If you are curious to know what the colors of the frames mean, you can click the Wireshark **View** menu and then click **Coloring Rules**. For example, ARP packets are light blue. You can also create your own color rules. Click the **–** button to the left of Frame to collapse the frame information.
5. Click the **+** button to the left of the Ethernet row, which is just below the Frame row you examined in Step 4. The section expands to display the Data Link layer header, which provides the source and destination MAC addresses. It also contains the protocol specification for the next layer. In this case, the packet must be passed up from the Data Link layer to ARP in the Network layer. Thus, 0806 specifies ARP as the next protocol. The "0x" indicates that the subsequent number, 0806, is hexadecimal. Click the **–** button to the left of Ethernet to collapse the Data Link information.
6. Click the **+** button to the left of the Address Resolution Protocol row, which is just below the Ethernet row. The section expands to display a variety of information; most importantly, this is an ARP request, as indicated by the fact that the sender's MAC and IP address are given but only the target's IP address is given. The sender does not know the target's MAC address, which is why the sender is making the ARP request. Click the **–** button to the left of Address Resolution Protocol to collapse that section.
7. Examine the rest of the ARP packets to be sure you understand what is happening in them. Use Table 2-16 to identify the different systems by their MAC and IP addresses.
8. Select one of the Echo (ping) reply packets in the top frame of the Wireshark window. Notice that an IP header and an ICMP data section now appear in the middle frame. Expand the Internet Protocol section and review each of the header fields, as explained earlier in this chapter.
9. Expand the Internet Control Message Protocol section. What is the pattern in the 32 bytes of data that are sent in a Windows ping? Leave the systems running for the next project.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Project 2-4: Examining IPv6 Ping Packets

Time Required: 30 minutes

Objective: Examine ND and ICMPv6 packets.

Description: In this project, you see the difference in how IPv6 handles the ping function compared with how IPv4 handled it in Hands-On Projects 2-2 and 2-3. You also see how Neighbor Discovery replaces ARP. Finally, you create a diagram of an ND header and explore IPv6 ping packets.

1. Access both the Windows 7 and Windows Server 2008 computers with administrative accounts.
2. Set up the desktop as you did in Step 4 of Hands-On Project 2-2, so that both the Wireshark window and the command prompt window are accessible.
3. As in Hands-On Project 2-2, ping your partner's system, but this time use the IPv6 address.
4. Your result should be similar to that shown in Figure 2-16. The first two packets are IPv6 Neighbor Discovery (ND) processes—specifically, Neighbor Solicitation messages. In Figure 2-16, the source IPv6 address is that of the Windows 7 computer. Search your capture for the equivalent packet. Notice that the destination address is not a unicast link-local address because it does not begin with fe80. The address starts with "ff," as do all IPv6 multicast addresses. This particular prefix, ff02::1, indicates a link-local scope all-nodes multicast address. (A good source for details like this is "Introduction to IPv6," which you can download at *http://technet.microsoft.com/en-us/library/bb726944.aspx*.) The middle frame should contain three headers: Ethernet, Internet Protocol, and Internet Control Message

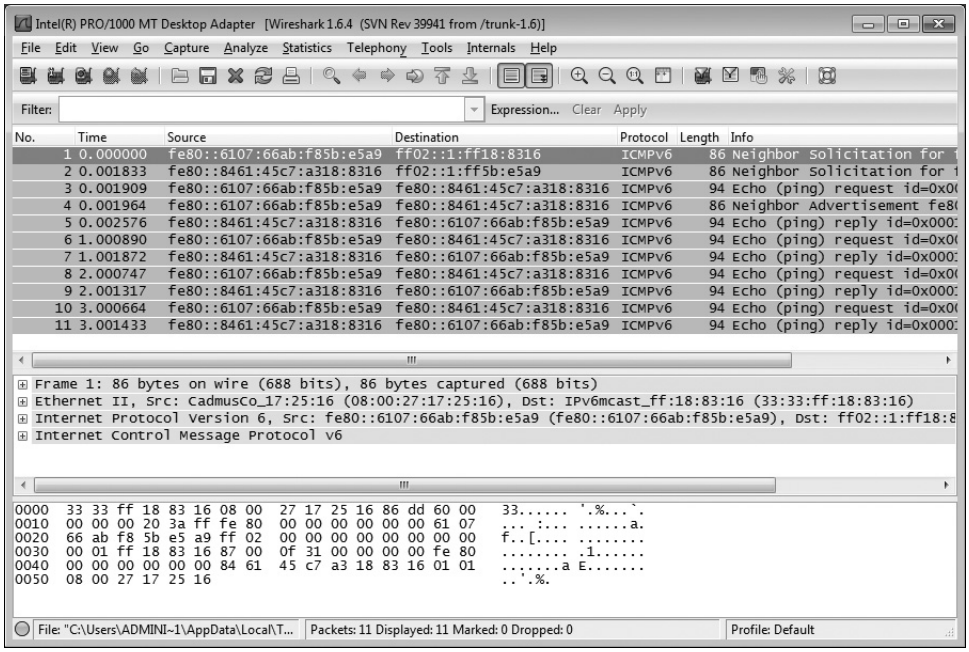

Source: Wireshark

Figure 2-16 Ping using IPv6 address

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Protocol. Investigate each of these headers in light of the information provided in this chapter. Pay particular attention to the Neighbor Solicitation message. Compare the Type field value to those in Table 2-15 to verify the type of ICMPv6.

5. This chapter had several figures that show header structures, including Figures 2-1, 2-3, 2-5, 2-8, and 2-9. In the following space, create a figure that shows the structure of the Neighbor Solicitation message.

6. Explore the Echo Request and Echo Reply packets in the rest of the capture. Verify the ICMPv6 message types by consulting Table 2-11.
7. Leave your system running for the next project.

Hands-On Project 2-5: A Challenge

Time Required: 20 minutes

Objective: Explore a ping option and independently analyze a packet capture.

Description: In this project, you experiment with a variation of the ping command and then analyze a packet capture to determine how the result is different from the one you obtained in Hands-On Project 2-4.

1. Access both the Windows 7 and Windows Server 2008 computers with administrative accounts.
2. Set up your system to capture a ping of your partner's system using the IPv6 address, as you have in the previous projects.
3. Enter the command **ping -l 5000** ***IPv6_address***, where -l is the letter *l*, not the number one, and *IPv6_address* is the IPv6 address of your partner's system. Capture the result.
4. Analyze the results. What is the effect of using the -l option with the ping command? How is this result different from that in Hands-On Project 2-4? How does IP handle this difference?
5. Log off your systems.

Case Projects

Case Project 2-1: The Differences Between IPv4 and IPv6

You are a network engineer for an IT consulting firm named F1IT. One of your clients, Beautivision, a chain of plastic surgery clinics with 80 locations nationwide, has asked you to prepare a proposal for implementing IPv6 in Beautivision's corporate headquarters, its WAN network, and its clinic locations. In preparation, you need to create a one- to two-page memo that describes the main differences between IPv4 and IPv6. Write the memo to Mary Jane Newman, communications manager for Beautivision. When preparing the memo, keep in mind that Ms. Newman has some knowledge of information technology but is not an expert.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 3

Network Traffic Signatures

After reading this chapter and completing the exercises, you will be able to:

- Explain the Common Vulnerabilities and Exposures (CVE) standard
- Describe the concepts of signature analysis
- Detect normal and suspicious traffic signatures
- Identify suspicious events

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

At its most basic level, securing network traffic can be broken down to a simple principle: You allow the communications that you want to pass through your network gateways, and you block the traffic you do not want. The challenge is to separate the two types of network traffic. To meet this challenge, you use a combination of firewalls, router access lists, intrusion detection and prevention system (IDPS) filtering, antivirus software, and other security tools. These devices determine whether traffic is normal or unauthorized based on criteria defined by system administrators. After a device can accurately tell the difference, it must then know how to respond, if it is capable.

This chapter begins by examining a standard for recording information about attack signatures: the Common Vulnerabilities and Exposures (CVE) standard. The CVE is a useful reference that helps you identify the signatures of current attacks.

This chapter examines techniques for identifying what constitutes normal network traffic and what constitutes an attempt to gain unauthorized access to network resources. You learn about signature analysis and how it is used in network security. You learn how to capture packets for inspection and analyze traffic signatures, both normal and suspicious. You then identify and analyze common suspicious events.

Examining the Common Vulnerabilities and Exposures Standard

One way to prevent attacks is to make sure your security devices can share information and coordinate with one another. At the perimeter of any network, you are likely to have a variety of hardware and software devices that provide security and that need to work cooperatively with one another. You might have a router from one vendor, a firewall from another, and an IDPS from a third. Unfortunately, the way these devices interpret signatures might differ. They probably address the same known attacks but name them differently and describe their characteristics differently. The **Common Vulnerabilities and Exposures** (**CVE**) standard enables these devices to share information about attack signatures and other vulnerabilities so that they can work together.

How CVE Works

CVE enables hardware and security devices that support it to draw from the same databases of vulnerabilities, which are presented in a standard format. For instance, a **scanner** is a device that scans a network for open ports or other potential vulnerabilities. If the scanner supports CVE, you can use it to compile a report that lists weak points in the system. When an alarm message is transmitted by an IDPS that also supports CVE, the attack signature can be compared to the report of current vulnerabilities to see whether an attack has actually occurred.

As you can see in Figure 3-1, the CVE standard affects many parts of a network:

1. A possible attack is detected by an IDPS sensor.
2. The signature is checked for a match against the database of known attack signatures available to the IDPS. If the IDPS being used is also CVE compliant, the report on the

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

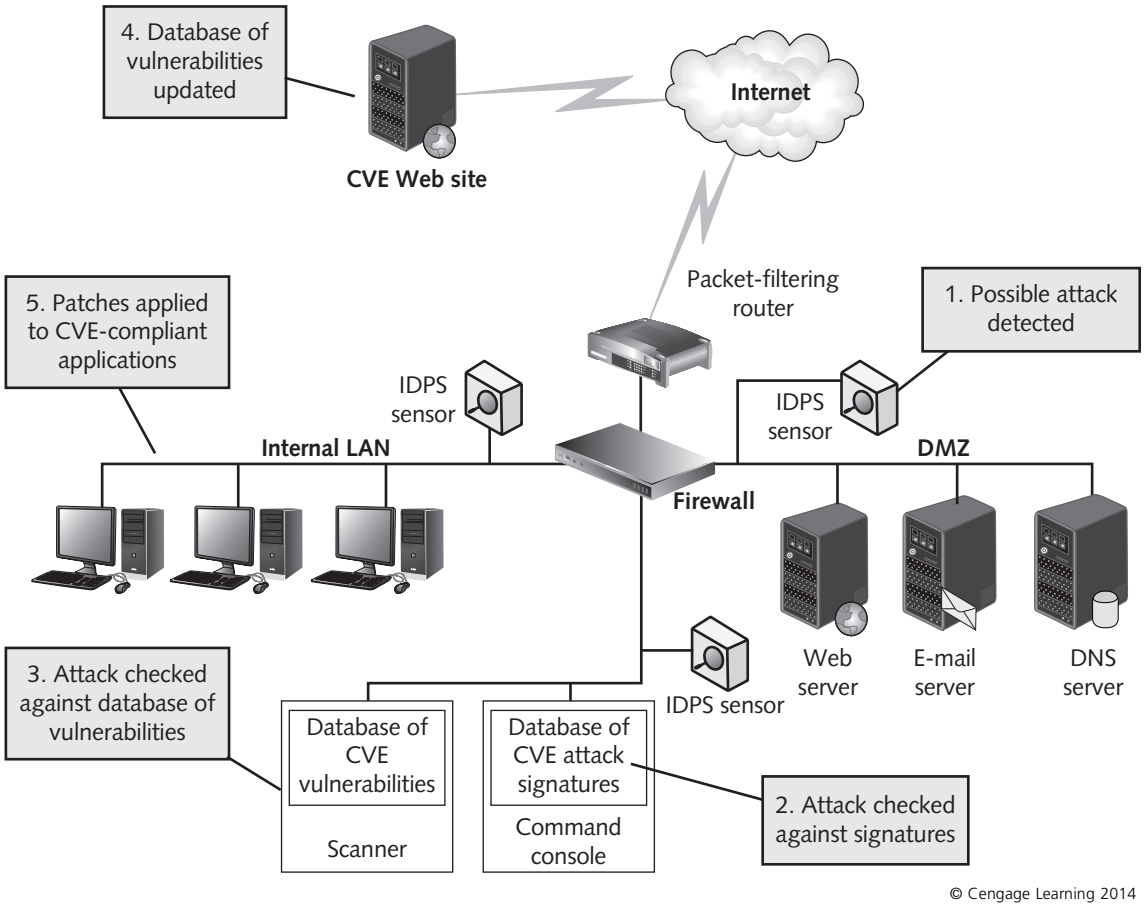

Figure 3-1 CVE enables multiple devices to work together to detect possible attacks

attack contains information on known network vulnerabilities associated with the attack signature.

3. The list of vulnerabilities is compared against a database of current vulnerable points in the system that have been compiled and stored by a CVE-compliant scanner to determine whether this possible attack can have an impact on the network.
4. Periodically, the list of vulnerabilities is updated with new entries from the CVE vulnerability Web site.
5. The manufacturers of CVE-compliant applications generate patches and updates in response to vulnerabilities. Those patches can then be applied to applications on the network.

Great benefits, such as stronger security and better performance, result when all the security devices on a network understand and use information that complies with the CVE standard. If you are responsible for purchasing an IDPS and other equipment for your organization, you should make sure they support CVE.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The CVE standard is a cooperative effort. Mitre maintains the database of vulnerabilities at *www.cve.mitre.org*. CVE has also served as the basis for new services. The National Institute of Standards and Technology, in a cooperative effort with the Department of Homeland Security, launched the Security Content Automation Protocol (SCAP) in 2009. SCAP serves as a repository of security checklists that provide low-level guidance on security configurations of operating systems and applications. SCAP is synchronized with the CVE list. You can learn more about SCAP at *http://scap.nist.gov/index.html*.

Scanning CVE Vulnerability Descriptions

You can go online to view current CVE vulnerabilities and even download the list so that you can review it at your convenience. However, as Mitre points out, the CVE list is not a vulnerability database that you can use to repair attacks on an IDPS; it is simply an informational tool. CVE listings are brief and simply refer to listings in other databases; they do not contain IP addresses, protocol listings, or other characteristics of an event that qualify it as a signature. A CVE reference reports the following information:

- The name of the vulnerability
- A short description
- References to the event in other databases, such as BUGTRAQ

The number associated with a CVE listing tells you when the listing was created. For instance, CVE-2012-0390 (shown in Figure 3-2) indicates that the listing was made in 2012 and was number 390 for that year.

CVE-ID

CVE-2012-0390 (under review)

Learn more at National Vulnerability Database (NVD)
• Severity Rating • Fix Information • Vulnerable Software Versions • SCAP Mappings

Description

The DTLS implementation in GnuTLS 3.0.10 and earlier executes certain error-handling code only if there is a specific relationship between a padding length and the ciphertext size, which makes it easier for remote attackers to recover partial plaintext via a timing side-channel attack, a related issue to CVE-2011-4108.

References

Note: References are provided for the convenience of the reader to help distinguish between vulnerabilities. The list is not intended to be complete.

- MISC:http://www.isg.rhul.ac.uk/~kp/dtls.pdf

Status

Candidate	This CVE Identifier has "Candidate" status and must be reviewed and accepted by the CVE Editorial Board before it can be updated to official "Entry" status on the CVE List. It may be modified or even rejected in the future.

Phase

Assigned (20120105)

Votes

Comments

Source: www.cve.mitre.org

Figure 3-2 CVE candidate listing CVE-2012-0390

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The listing in Figure 3-2 indicates that the vulnerability is associated with the Gnu Datagram Transport Layer Security (TLS) protocol, which is associated with open-source Web security applications. The reported vulnerability is under review as a CVE candidate.

A candidate for inclusion in CVE must be reviewed and accepted by the CVE Editorial Board before it can be added to the CVE list. The CVE numbering system includes a status line indicating whether the vulnerability has a status of Candidate, Entry, or Deprecated. Only the status line of a CVE is updated when new CVE versions are released. You can view vulnerability listings like the one in Figure 3-2 at *www.cve.mitre.org*. Use the CVE's search engine to look for specific vulnerabilities. The entire list is very long, so search for a specific year or keyword, such as that of an application, file, or device, to produce a smaller list.

Understanding Signature Analysis

A **signature** is a set of characteristics—such as IP numbers and options, TCP flags, and port numbers—that define a type of network activity. Besides individual TCP/IP packet attributes, a signature can also consist of a sequence of packets or other events, such as logons to a network.

Some intrusion-detection devices assemble databases of "normal" traffic signatures. As traffic is detected, it is compared to the database, and any deviations from normal signatures (the network baseline) trigger an alarm. Intrusion-prevention devices can also go further and drop the packets, and perhaps create firewall filtering rules to prevent similar traffic from entering the network. Other intrusion-detection devices refer to a database of well-known attack signatures. Any traffic that matches one of the stored attack signatures triggers an alarm or causes the IDPS to take steps to prevent the attack. Your understanding of normal and suspicious traffic signatures enables you to configure an IDPS to work more effectively by minimizing the number of false positives (false alarms) and decreasing the number of false negatives (missed attacks). You learn about false positives and false negatives in Chapter 8.

The following sections introduce you to signature analysis. Then, you learn more about analyzing packets and review normal traffic signatures you are likely to encounter. Finally, you learn about suspicious traffic signatures that indicate a possible attempt to scan and gain unauthorized access to your network.

Signature analysis is the practice of analyzing and understanding TCP/IP communications to determine whether they are legitimate or suspicious. Packets are the most basic level of network communications.

Suspicious TCP/IP packets fall into several categories: bad header information, suspicious data payload, single-packet attacks, and multiple-packet attacks.

Bad Header Information

Packets are commonly altered through their header information, and packet filters usually scan for these alterations. Suspicious signatures can include malformed data that affect some or all of the following:

- Source and destination IP address
- Source and destination port number
- IP options

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- IP fragmentation flags, fragmentation offset, or fragment identification
- IP protocol
- IP, TCP, or UDP checksums

A **checksum** is a simple error-checking procedure for determining whether a message has been tampered with or damaged in transit. A mathematical formula is used to process the number of data bits in a message. A numeric value (the checksum) is then calculated. The receiving computer applies the same formula to the message; if a different checksum is found, the receiving computer determines that the message has been tampered with or corrupted and drops it.

Attackers can use software that generates packets set to their specifications to forge IP addresses or other types of header information. For instance, a packet can be broken into "chunks" and sent in a series. An attacker can eliminate the initial chunk in the series from the set, which makes the receiving computer unable to reassemble the packets. Therefore, the series of chunks that follow can circumvent a packet filter. An attacker can send more or fewer packets than indicated in the initial packet, which could disable a server that cannot process a different number of packets than it expected to receive.

Suspicious Data Payload

The payload (data) part of a packet is the actual data sent from an application on one computer to an application on another. Sometimes, attacks can be detected by an IDPS that matches a text string to a specific set of characters in the payload. For instance, a program called CyberEYE creates **remote-access Trojans (RATs)**; when installed on unsuspecting systems, these RATs open back doors that give the remote attacker administrative rights on the victim's computer. The CyberEYE program was created in Turkey, and one of its signatures is the presence of the following hexadecimal string in the data portion of the packet: 41 4E 41 42 49 4C 47 49 7C. When converted into ASCII, this string becomes ANA BILGI, which is Turkish for MAIN INFORMATION, according to security analyst Chris Sanders. The presence of this hexadecimal string in the data portion of the packet is a key part of the CyberEYE session establishment detection signature, particularly when the packet is communicated over TCP port 4433, a port frequently used by CyberEYE.

In another type of attack, the UNIX Sendmail program is exploited by adding codes to packet contents. Codes such as VRFY and EXPN are used to uncover account names on the Sendmail server. By adding the code EXPN DECODE in a packet's data payload, attackers attempt to establish a connection with an alias called "decode." If a connection is made, attackers can use it to place malicious files on the exploited system. To defend against this type of attack, a network administrator should remove the "decode" alias line, which is installed by default with many UNIX/Linux systems in the /etc/mail/aliases file.

Internet searches can yield a wealth of tools and information, but be careful when downloading files. You might end up downloading the malicious software you are trying to avoid! Create a folder to store the downloaded files, unzip the contents into the new folder, if necessary, and then run a virus scan on the folder. Any known malware signature will be recognized, provided that you keep your antivirus software updated, and you can safely remove the files without harm in most cases. You can reduce risk by using common sense and basic security protocols.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Single-Packet Attacks

A **single-packet attack** (also called an "atomic attack") can be completed by sending a single network packet from a client to a host. Because only one packet is needed, a connection does not need to be established between the two computers. Many changes to IP option settings can cause a server to freeze up because it does not know how to handle these packets. The IP option settings are shown in Table 3-1.

Option number	Name of option
0	End of Options
1	No Operation
2	Security
3	Loose Source and Record Routing
4	Internet Timestamp
7	Record Return Route
8	Option has been deprecated
9	Strict Source and Record Routing

© Cengage Learning 2014

Table 3-1 IP option settings

As an example of IP options processing, suppose an ICMP echo request (or "ping") packet is sent from a host to a server with Option 7 set. The echo reply response from the server might spell out the route the request takes to return from the server, thus revealing the IP addresses of hosts or routers on the network that the attacker can then target. Option 4 can be used with Option 7 to record the amount of time the echo reply packet spends between "hops" on the network. A hop is the movement of a packet from one point on a network to another.

Multiple-Packet Attacks

In contrast to single-packet attacks, **multiple-packet attacks** (also called "composite attacks") require a series of packets to be received and executed. These attacks are especially difficult to detect. They require an IDPS to have multiple attack signatures on hand for reference. In addition, the IDPS sensor needs to maintain state information about a connection after it has been established, and it needs to keep that state information on hand for the entire length of an attack.

Denial of service (DoS) attacks are obvious examples of multiple-packet attacks. A type of DoS attack called an ICMP flood occurs when multiple ICMP packets are sent to a single host on a network. The result of this flood is that the server becomes so busy responding to the ICMP requests that it cannot process other traffic.

Analyzing Packets

In Chapter 2, you worked with a packet sniffer called Wireshark. A **packet sniffer** captures information about each TCP/IP packet it detects. By using such an application, you can

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

study packets and identify characteristic features that tell you what type of connection is under way and whether the transmission is legitimate or suspicious.

Capturing packets and studying them can help you better understand what makes up a signature. Figure 3-3 shows a type of packet you have already seen: an ICMP ping packet. Now examine this ICMP header more closely.

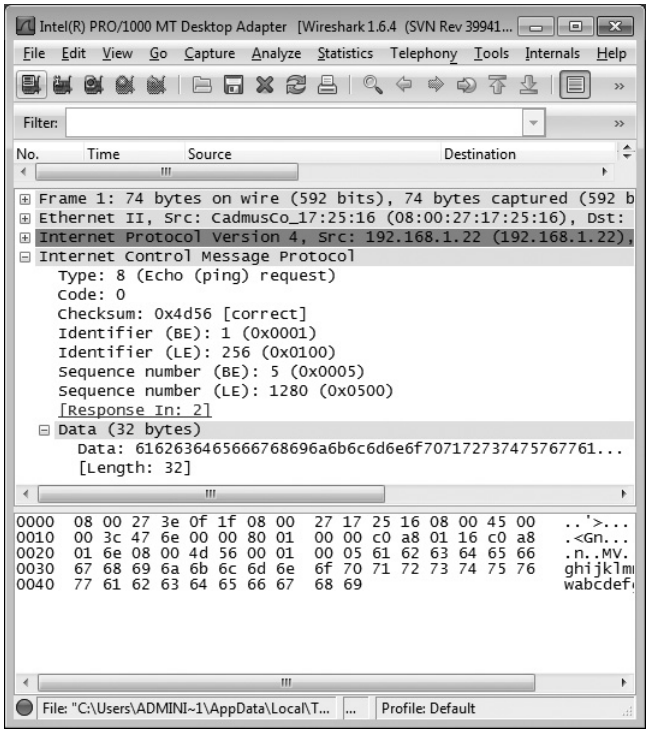

Source: Wireshark

Figure 3-3 An ICMP echo request packet capture

This capture was made using Wireshark. Packets are color-coded for various protocols and transmission errors; for example, checksum errors are shown in black with a red font. ICMP packets are shown in blue with a black font. You can customize the color coding to follow conversations, or open individual packets or conversations in a new window.

The following list describes the elements in the ICMP header:

- *Type*—ICMP has different types of messages, such as echo request, redirect, and source quench. Type 8 is an echo request packet. (See Table 2-7 for a list of ICMP types.)
- *Code*—This 8-bit value provides further information about some types of ICMP packets.
- *Checksum*—This 16-bit field contains error-checking data.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *Identifier (BE)*—This field and the following field contain **ICMP ID numbers**. The (BE) stands for Big Endian, which refers to how an operating system stores this ICMP information in memory. Big Endian stores the most significant byte first.
- *Identifier (LE)*—This field also contains an ICMP ID number. The (LE) stands for Little Endian, which stores the least significant byte first.
- *Sequence number (BE)*—This field reports the ICMP sequence number used to identify the ICMP packet in a sequence of packets.
- *Sequence number (LE)*—This field also reports the ICMP sequence number used to identify the ICMP packet in a sequence of packets.
- *Response In*—This field indicates receipt of a response to this request. This field also indicates the packet number of the response—in this case, number 2.
- *Data*—This field contains the actual data the packet is communicating.

You examined a TCP header in Chapter 2, but you now look at the header structure in more detail. Figure 3-4 shows a TCP header.

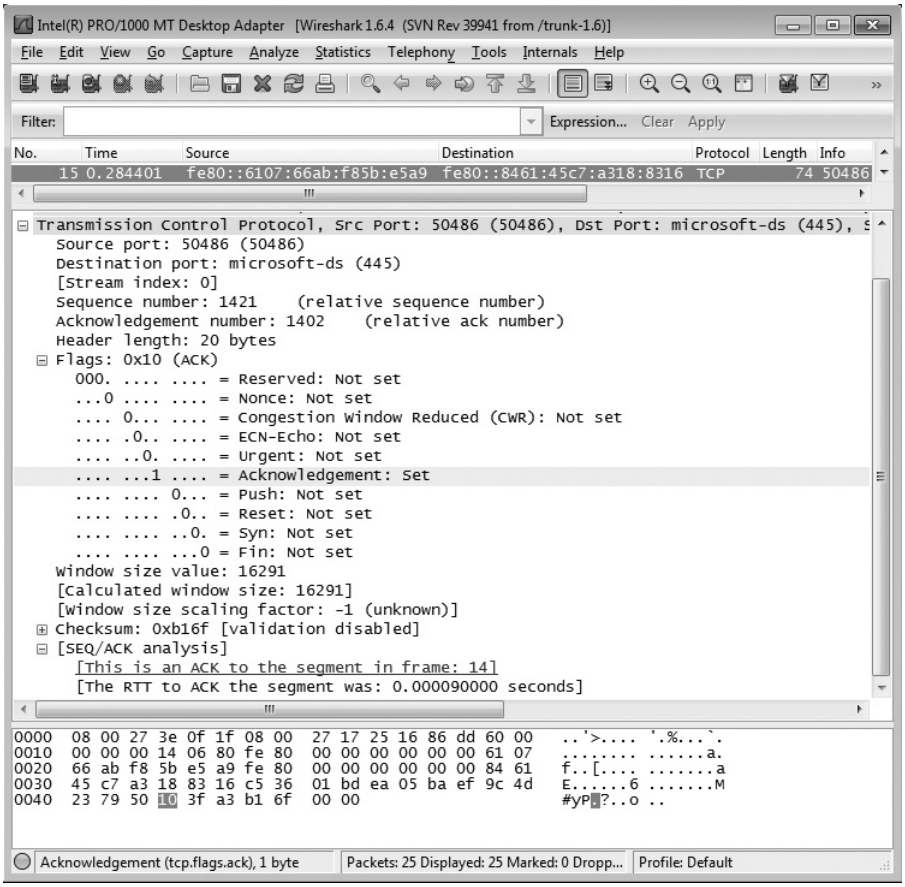

Source: Wireshark

Figure 3-4 A TCP packet capture

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The following list describes the TCP-specific elements in the packet:

- *Source port*—This field shows the source port of 50486. This is a high port within the range that a host would assign itself dynamically.
- *Destination port*—This field shows the destination port of 445. Wireshark has identified the port as "microsoft-ds," which stands for Microsoft Domain Services. In this case, the packet is associated with Windows 7 accessing a shared folder on Windows Server 2008. Originally, Microsoft used NetBIOS ports (135–139) for file sharing; however, after serious security problems occurred, Microsoft switched to Server Message Blocks over TCP and used port 445. Note that SMB has been replaced by CIFS (Common Internet File System), but most utilities like Wireshark still identify this traffic as SMB.
- *Stream index*—The stream index is an internal Wireshark mapping that identifies a unique TCP stream. All packets for the same TCP stream should have the same values in this field, although source and destination addresses are switched for two-way communications.
- *Sequence number*—This value provides the sequence number of the packet.
- *Acknowledgement number*—This number acknowledges receipt of the previous packet in the sequence. Remember that this number is one higher than the sequence number in the packet that the sender just received.
- *Header length*—This value is the overall length of the TCP header plus options, if there are any.
- *Flags*—In the packet shown in Figure 3-4, only the ACK flag is set. The ACK (acknowledgement) flag indicates that a connection has been established.
- *Window size value*—This value indicates the size of the window (buffer size) on the source computer so that the recipient can determine how many packets can be sent at any one time.
- *Checksum*—This 16-bit field contains error-checking data.
- *SEQ/ACK analysis*—This value indicates that the packet is an ACK response to another frame. The value is provided by Wireshark and is not part of the actual header. If the packet is the first in a TCP session, this field does not appear.

You can use a variety of software tools to capture packets that pass through a computer's network interface card. For example, Wireshark has a command-line utility called TShark, which you use in the next section. Some traffic patterns are easier to explore from the command line.

Analyzing Traffic Signatures

Now that you have learned the basics of packet captures, you must learn to determine whether traffic is normal or suspicious. You are probably familiar with the concept of network baselining. You must first know what is normal for your network before you can identify anomalies. This section helps you learn to tell the difference between normal traffic and suspicious activity.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Examining Normal Network Traffic Signatures

You have looked at normal packets already, but now you look for specific normal patterns in groups of packets. Wireshark includes a command-line tool for packet captures called TShark, which you use in the Hands-On Projects at the end of this chapter. In Figure 3-5, you can see a ping sequence captured by TShark. In this capture, it is easier to see the packet flow in the abbreviated command-line format.

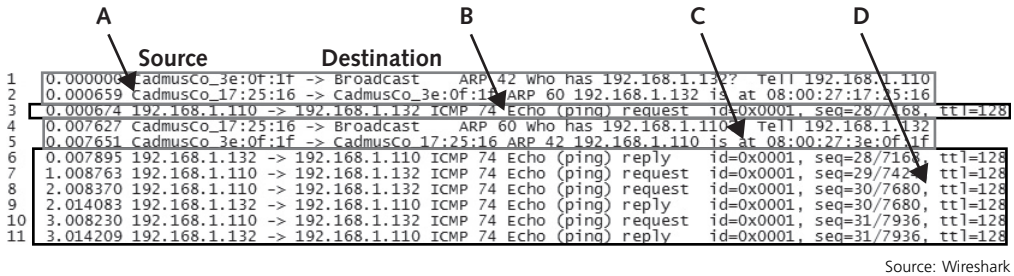

Source: Wireshark

Figure 3-5 Abbreviated ping capture with TShark

In Figure 3-5, the box marked "A" shows the first packets sent after the Windows 7 system at 192.168.1.132 pinged 192.168.1.110, the IP address of the Windows Server 2008 computer. As you might recall, Windows 7 must send an ARP query to determine the MAC address of the Windows Server 2008 computer. This activity is completed in packets 1 and 2. In packet 3, in the box marked "B," Windows 7 sends the first echo request to Windows Server 2008. In box "C," Windows Server 2008 initiates the ARP query to determine the Windows 7 MAC address. Finally, in the packets in box "D," the remainder of the echo requests and replies are completed.

One aspect of normal TCP signatures that is easiest to identify is the use of flags. TCP flags were presented in Chapter 2; the most important flags are described here in more detail.

- *SYN (synchronize) flag (0x2)*—This flag is sent from one computer to another when a connection is initiated; the two computers are attempting to synchronize a connection. In the TCP three-way handshake, the initiator of the communication sends a packet with the SYN flag set. The normal response is a TCP packet with the SYN and ACK flags set. The initiator then responds with an ACK flag set.
- *ACK (acknowledgement) flag (0x10)*—The ACK flag is set to signal that the previous packet was received. This number is one higher than the sequence number of the packet just received. The sequence and acknowledgement numbers are what makes TCP a connection-oriented protocol.
- *PSH (push) flag (0x8)*—This flag indicates that immediate data delivery is required and to forward all of the queued data (data stored in a buffer) to the destination protocol layer immediately.
- *URG (urgent) flag (0x20)*—This flag is used when urgent data is being sent from one computer to another.
- *RST (reset) flag (0x4)*—This flag is sent when one computer wants to stop and restart the connection in response to a problem. On some systems, this flag is the default response to a probe on a closed port.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *FIN (finished) flag (0x1)*—This flag lets one computer know that the other is finished sending data. This flag might also be a response to a probe on a closed port.

The placement and use of these flags are strictly defined, and deviations from normal use mean that the communication is suspicious. For instance, the SYN flag should appear at the beginning of a connection; the FIN flag should appear only at the end of a connection. The presence of both the SYN and FIN flags in the same packet makes no logical sense and is not a "legal" TCP configuration. This packet would definitely indicate suspicious network activity. However, the ACK and PSH flags can be used together when data is sent from one computer to another.

Figure 3-6 shows an abbreviated capture of a stream of TCP traffic. (Note that TShark does not normally list the packet number; these have been added for clarity.)

	Seconds from start of capture	Source IP address	Destination IP address	Protocol	Size	Source Port	Destination Port	TCP flags	Sequence/Acknowledgement numbers, etc.
1	18.743157	192.168.1.132	192.168.1.1	TCP	66	49263	1780	[SYN]	Seq=0 Win=8192 Len=0 MSS=1460 WS=4 SACK_PERM=1
2	18.748737	192.168.1.1	192.168.1.132	TCP	66	1780	49263	[SYN, ACK]	Seq=0 Ack=1 Win=5840 Len=0 MSS=1460 SACK_PERM=1 WS=2
3	18.748920	192.168.1.132	192.168.1.1	TCP	54	49263	1780	[ACK]	Seq=1 Ack=1 Win=65700 Len=0
4	18.749473	192.168.1.132	192.168.1.1	TCP	243	49263	1780	[PSH, ACK]	Seq=1 Ack=1 Win=65700 Len=189
5	18.758879	192.168.1.1	192.168.1.132	TCP	60	1780	49263	[ACK]	Seq=1 Ack=190 Win=5840 Len=0
6	18.759032	192.168.1.1	192.168.1.132	TCP	156	1780	49263	[PSH, ACK]	Seq=1 Ack=190 Win=5840 Len=102
7	18.760956	192.168.1.1	192.168.1.132	TCP	1514	1780	49263	[ACK]	Seq=103 Ack=190 Win=5840 Len=1460
8	18.761105	192.168.1.132	192.168.1.1	TCP	54	49263	1780	[ACK]	Seq=190 Ack=1563 Win=65700 Len=0
9	18.775169	192.168.1.1	192.168.1.132	TCP	1514	1780	49263	[ACK]	Seq=1563 Ack=190 Win=5840 Len=1460
10	18.777046	192.168.1.1	192.168.1.132	TCP	1514	1780	49263	[ACK]	Seq=3023 Ack=190 Win=5840 Len=1460
11	18.777213	192.168.1.132	192.168.1.1	TCP	54	49263	1780	[ACK]	Seq=190 Ack=4483 Win=65700 Len=0
12	18.777292	192.168.1.1	192.168.1.132	TCP	122	1780	49263	[FIN, PSH, ACK]	Seq=4483 Ack=190 Win=5840 Len=68
13	18.777446	192.168.1.132	192.168.1.1	TCP	54	49263	1780	[ACK]	Seq=190 Ack=4552 Win=65632 Len=0
14	18.778399	192.168.1.132	192.168.1.1	TCP	54	49263	1780	[FIN, ACK]	Seq=190 Ack=4552 Win=65632 Len=0
15	18.783569	192.168.1.1	192.168.1.132	TCP	60	1780	49263	[ACK]	Seq=4552 Ack=191 Win=5840 Len=0

Source: Wireshark

Figure 3-6 TShark capture of a TCP stream

A packet-by-packet analysis reveals:

- *Packet 1*—The host at 192.168.1.132 has sent a SYN packet to the host at 192.168.1.1. The sending host is using port 49263 and is sending to port 1780 of the destination system. TShark assigns relative sequence and acknowledgement numbers in its captures, so it is not a coincidence that the initial sequence number in this packet is 0. If you looked at the actual data in the capture, you would see a much more random and larger sequence number. You would also see a 0 as the acknowledgement number for this packet, but because this is the first packet in the communication session, Wireshark omits the acknowledgement number in this abbreviated capture report. Note also that the sliding window size (Win) is 8192, the maximum segment size (MSS) is indicated, and the sending host is capable of using selective acknowledgements (SACK_PERM=1). **Selective acknowledgements** speed up transmissions by allowing a receiver to send acknowledgements for specific packets so that the sender can resend only those packets that were lost.
- *Packet 2*—The receiver of the first packet at 192.168.1.1 completes the second stage of the TCP three-way handshake by sending a SYN ACK packet from its TCP port 1780 to the host at port 49263 of 192.168.1.132. Note that all the packets in this session use these same ports. While changes in active ports during a session can be

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

normal, such activity can sometimes be suspicious. Sliding window size, SACK capability, and other items are negotiated.

- *Packet 3*—The TCP three-way handshake is completed when the initiator of the session sends an ACK packet.
- *Packets 4–11*—These packets transport data. Notice that the size of the first three packets was not greater than 66 bytes. Packets 4–11 are much larger, and might contain up to 1514 bytes—the maximum for Ethernet packets. Notice also that each of these packets contains an ACK flag or a PSH and ACK flag. Notice too that packets 5, 6, 7, 9, and 10 are sent by the host at 192.168.1.1. Because of the sliding windows, the host at 192.168.1.132 does not have to send an acknowledgement for each packet one at a time. In packets 7, 9, and 10, it appears that the data is being sent because all three packets are from 192.168.1.1 and their size is 1514 bytes. On the other hand, some of the ACK packets from 192.168.1.132 are only 54 bytes, which presumably indicates a simple acknowledgement that packets were received. These small packets are not sending payload data.
- *Packet 12*—The system that seems to be sending the most data—perhaps a server, at 192.168.1.1—sends a FIN, PSH, ACK packet. This indicates that the server has completed the requested data transfer.
- *Packet 13*—192.168.1.132 acknowledges the FIN.
- *Packet 14*—192.168.1.132 sends its own FIN ACK.
- *Packet 15*—The session is closed when 192.168.1.1 acknowledges the FIN from 192.168.1.132.

Note the sequence for tearing down a TCP session:

Host A: FIN

Host B: ACK

Host B: FIN ACK

Host A: ACK

FTP Signatures If your organization operates a public FTP server, you will be called on regularly to review the signatures of packets that attempt to access that server. You need to determine whether the computer that makes the connection attempt is allowed to access the server in accordance with your packet-filtering rules.

The signature of a normal connection between a client and an FTP server includes a three-way handshake. The first three packets shown in Figure 3-7 contain different TCP flags that enable you to keep track of the connection.

You can analyze the packet sequence shown in Figure 3-7 as follows:

1. In the first packet, the computer at IP address 192.168.1.132:50580 attempts to connect to the FTP server at 192.168.1.110:21.

The syntax *IP address:port* is called a socket. In this packet, the computer at 192.168.1.132 is using its port 50580, and is sending to port 21 of the server at 192.168.1.110.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

```
| 0.000000 192.168.1.132 -> 192.168.1.110 TCP 66 50580 > 21 [SYN] Seq=0 Win=8192 Len=0 MSS=1460

0000  08 00 27 3e 0f 1f 08 00 27 17 25 16 08 00 45 00   ..'>....'.%...E.
0010  00 34 19 de 40 00 80 06 00 00 c0 a8 01 84 c0 a8   .4..@...........
0020  01 6e c5 94 00 15 54 9d 21 a7 00 00 00 00 80 02   .n....T.!.......
0030  20 00 84 69 00 00 02 04 05 b4 01 03 03 00 01 01    ..i............
0040  04 02                                            ..

  0.000533 192.168.1.110 -> 192.168.1.132 TCP 66 21 > 50580 [SYN, ACK] Seq=0 Ack=1 Win=8192 Len

0000  08 00 27 17 25 16 08 00 27 3e 0f 1f 08 00 45 00   ..'.%...'>....E.
0010  00 34 24 af 40 00 80 06 51 d2 c0 a8 01 6e c0 a8   .4$.@...Q....n..
0020  01 84 00 15 c5 94 44 6e f5 cb 54 9d 21 a8 80 12   ......Dn..T.!...
0030  20 00 54 94 00 00 02 04 05 b4 01 03 03 08 01 01    .T.............
0040  04 02                                            ..

  0.000648 192.168.1.132 -> 192.168.1.110 TCP 54 50580 > 21 [ACK] Seq=1 Ack=1 Win=8192 Len=0

0000  08 00 27 3e 0f 1f 08 00 27 17 25 16 08 00 45 00   ..'>....'.%...E.
0010  00 28 19 df 40 00 80 06 00 00 c0 a8 01 84 c0 a8   .(..@...........
0020  01 6e c5 94 00 15 54 9d 21 a8 44 6e f5 cc 50 10   .n....T.!.Dn..P.
0030  20 00 84 5d 00 00                                 ...]..

  0.001455 192.168.1.110 -> 192.168.1.132 FTP 81 Response: 220 Microsoft FTP Service

0000  08 00 27 17 25 16 08 00 27 3e 0f 1f 08 00 45 00   ..'.%...'>....E.
0010  00 43 24 b0 40 00 80 06 51 c2 c0 a8 01 6e c0 a8   .C$.@...Q....n..
0020  01 84 00 15 c5 94 44 6e f5 cc 54 9d 21 a8 50 18   ......Dn..T.!.P.
0030  01 00 19 5b 00 00 32 32 30 20 4d 69 63 72 6f 73   ...[..220 Micros
0040  6f 66 74 20 46 54 50 20 53 65 72 76 69 63 65 0d   oft FTP Service.
0050  0a                                               .

  0.208118 192.168.1.132 -> 192.168.1.110 TCP 54 50580 > 21 [ACK] Seq=1 Ack=28 Win=8165 Len=0

0000  08 00 27 3e 0f 1f 08 00 27 17 25 16 08 00 45 00   ..'>....'.%...E.
0010  00 28 19 e0 40 00 80 06 00 00 c0 a8 01 84 c0 a8   .(..@...........
0020  01 6e c5 94 00 15 54 9d 21 a8 44 6e f5 e7 50 10   .n....T.!.Dn..P.
0030  1f e5 84 5d 00 00                                 ...]..
```

Source: Wireshark

Figure 3-7 The beginning of an FTP session

The packet has the SYN flag set because a synchronization request is being made to the remote server. The sequence number of the packet is set at 0 with a window size of 8192. There is no acknowledgement number because this is the first packet of the session; in other words, there is no previous packet to acknowledge.

2. In the second packet, the FTP server responds to the client by sending a packet from its port 21 with the ACK and SYN flags set. The server uses the same initial sequence number as the client used in packet 1, but the server increments the client's initial sequence number by one to create an acknowledgement number. This number is sent back to the client to acknowledge that the server received the first packet.
3. In the third packet, the client responds with a packet that has the ACK flag set. This concludes the three-way handshake.
4. In the fourth packet, the FTP server indicates that it is ready for a new user. This FTP packet is the only non-TCP packet shown in the figure.
5. In the fifth packet, the client acknowledges the server's last packet.

In Figure 3-8, you see the continuation of the FTP session.

1. The user enters the username. Note that this text appears in the ASCII portion of the packet.
2. In the second packet, the FTP server challenges the client for the user's password.
3. In the third packet, the client sends an acknowledgement.
4. In the fourth packet, the user's password traverses the network unencrypted. Anyone with a packet sniffer can learn that the password is Pa$$word.
5. In the fifth packet, the user learns that he is logged in successfully.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

```
  2.123399 192.168.1.132 -> 192.168.1.110 FTP 65 Request: USER dean

0000  08 00 27 3e 0f 1f 08 00 27 17 25 16 08 00 45 00   ..'>....'.%...E.
0010  00 33 19 e1 40 00 80 06 00 00 c0 a8 01 84 c0 a8   .3..@...........
0020  01 6e c5 94 00 15 54 9d 21 a8 44 6e f5 e7 50 18   .n....T.!.Dn..P.
0030  1f e5 84 68 00 00 55 53 45 52 20 64 65 61 6e 0d   ...h..USER dean.
0040  0a                                                .

  2.124041 192.168.1.110 -> 192.168.1.132 FTP 87 Response: 331 Password required for dean.

0000  08 00 27 17 25 16 08 00 27 3e 0f 1f 08 00 45 00   ..'.%...'>....E.
0010  00 49 24 b1 40 00 80 06 51 bb c0 a8 01 6e c0 a8   .I$.@...Q....n..
0020  01 84 00 15 c5 94 44 6e f5 e7 54 9d 21 b3 50 18   ......Dn..T.!.P.
0030  01 00 f9 d8 00 00 33 33 31 20 50 61 73 73 77 6f   ......331 Passwo
0040  72 64 20 72 65 71 75 69 72 65 64 20 66 6f 72 20   rd required for
0050  64 65 61 6e 2e 0d 0a                              dean...

  2.317076 192.168.1.132 -> 192.168.1.110 TCP 54 50580 > 21 [ACK] Seq=12 Ack=61 Win=8132 Le

0000  08 00 27 3e 0f 1f 08 00 27 17 25 16 08 00 45 00   ..'>....'.%...E.
0010  00 28 19 e2 40 00 80 06 00 00 c0 a8 01 84 c0 a8   .(..@...........
0020  01 6e c5 94 00 15 54 9d 21 b3 44 6e f6 08 50 10   .n....T.!.Dn..P.
0030  1f c4 84 5d 00 00                                 ...]..

  6.426319 192.168.1.132 -> 192.168.1.110 FTP 69 Request: PASS Pa$$word

0000  08 00 27 3e 0f 1f 08 00 27 17 25 16 08 00 45 00   ..'>....'.%...E.
0010  00 37 19 e3 40 00 80 06 00 00 c0 a8 01 84 c0 a8   .7..@...........
0020  01 6e c5 94 00 15 54 9d 21 b3 44 6e f6 08 50 18   .n....T.!.Dn..P.
0030  1f c4 84 6c 00 00 50 41 53 53 20 50 61 24 24 77   ...l..PASS Pa$$w
0040  6f 72 64 0d 0a                                   ord..

  6.427227 192.168.1.110 -> 192.168.1.132 FTP 75 Response: 230 User logged in.

0000  08 00 27 17 25 16 08 00 27 3e 0f 1f 08 00 45 00   ..'.%...'>....E.
0010  00 3d 24 b2 40 00 80 06 51 c6 c0 a8 01 6e c0 a8   .=$.@...Q....n..
0020  01 84 00 15 c5 94 44 6e f6 08 54 9d 21 c2 50 18   ......Dn..T.!.P.
0030  01 00 99 e6 00 00 32 33 30 20 55 73 65 72 20 6c   ......230 User l
0040  6f 67 67 65 64 20 69 6e 2e 0d 0a                  ogged in...
```

Source: Wireshark

Figure 3-8 Continuation of an FTP session

It is disturbing that FTP sends content as plain text. Evidently, this FTP server requires users to authenticate, yet there is no protection from interception. It would be much safer to allow only anonymous access in this case so that users' credentials would not be at risk. The administrator would lose the ability to monitor who accessed the server, but at least users' network passwords would be safe. An even better approach would be to encrypt FTP traffic.

The FTP control traffic, which includes authentication and directory browsing information, typically uses the FTP server's port 21. File transfer typically uses the FTP server's port 20.

In the last part of the FTP session (see Figure 3-9), you see the teardown of the FTP session with the use of the FIN flag.

1. In the first packet, the FTP server announces that it has completed the transfer by sending a FIN ACK packet.
2. In the second packet, the client acknowledges the FIN ACK.
3. In the third packet, the client sends its own FIN ACK packet to announce that it has completed the session as well.
4. The final packet of the session is the server's ACK of the previous FIN.

Web Signatures Most of the signatures you see in the log files you analyze will probably be Web-related. When a signature is Web-related, it consists of packets sent back and forth from a Web browser to a Web server as a connection is made. The signature of a

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

```
 15.632892 192.168.1.110 -> 192.168.1.132 TCP 60 21 > 50580 [FIN, ACK] Seq=312 Ack=113 '

0000  08 00 27 17 25 16 08 00 27 3e 0f 1f 08 00 45 00   ..'.%...'>....E.
0010  00 28 24 c4 40 00 80 06 51 c9 c0 a8 01 6e c0 a8   .($.@...Q....n..
0020  01 84 00 15 c5 94 44 6e f7 03 54 9d 22 18 50 11   ......Dn..T.".P.
0030  01 00 b2 bf 00 00 00 00 00 00 00 00               ............

 15.632960 192.168.1.132 -> 192.168.1.110 TCP 54 50580 > 21 [ACK] Seq=113 Ack=313 Win=7

0000  08 00 27 3e 0f 1f 08 00 27 17 25 16 08 00 45 00   ..'>....'.%...E.
0010  00 28 19 f4 40 00 80 06 00 00 c0 a8 01 84 c0 a8   .(..@...........
0020  01 6e c5 94 00 15 54 9d 22 18 44 6e f7 04 50 10   .n....T.".Dn..P.
0030  1e c9 84 5d 00 00                                 ...]..

 15.635862 192.168.1.132 -> 192.168.1.110 TCP 54 50580 > 21 [FIN, ACK] Seq=113 Ack=313 '

0000  08 00 27 3e 0f 1f 08 00 27 17 25 16 08 00 45 00   ..'>....'.%...E.
0010  00 28 19 f5 40 00 80 06 00 00 c0 a8 01 84 c0 a8   .(..@...........
0020  01 6e c5 94 00 15 54 9d 22 18 44 6e f7 04 50 11   .n....T.".Dn..P.
0030  1e c9 84 5d 00 00                                 ...]..

 15.636335 192.168.1.110 -> 192.168.1.132 TCP 60 21 > 50580 [ACK] Seq=313 Ack=114 Win=6

0000  08 00 27 17 25 16 08 00 27 3e 0f 1f 08 00 45 00   ..'.%...'>....E.
0010  00 28 24 c5 40 00 80 06 51 c8 c0 a8 01 6e c0 a8   .($.@...Q....n..
0020  01 84 00 15 c5 94 44 6e f7 04 54 9d 22 19 50 10   ......Dn..T.".P.
0030  01 00 b2 be 00 00 00 00 00 00 00 00               ............
```

Source: Wireshark

Figure 3-9 The teardown of an FTP data connection

normal handshake between two Web browsers consists of a sequence of packets that are distinguished by their TCP flags.

As mentioned in the preceding section, normal TCP traffic uses several TCP flags to control the connection. Being aware of these flags and where they are used can help you determine whether a signature is normal or part of a possible intrusion attempt.

In Figure 3-10, five packets are associated with the handshake between the Web browser at IP address 192.168.1.132:50969 and the Web server 65.54.87.10:80 (the typical HTTP port used for Web traffic).

```
  1.825108 192.168.1.132 -> 65.54.87.10  TCP 66 50969 > 80 [SYN] Seq=0 Win=8192 Len=0 MSS=1460 WS=4 :

0000  58 6d 8f 04 8a d3 08 00 27 17 25 16 08 00 45 00   Xm......'.%...E.
0010  00 34 25 71 40 00 80 06 00 00 c0 a8 01 84 41 36   .4%q@.........A6
0020  57 0a c7 19 00 50 1a f0 b1 44 00 00 00 00 80 02   W....P...D......
0030  20 00 5a 93 00 00 02 04 05 b4 01 03 03 02 01 01    .Z.............
0040  04 02                                             ..

  1.849810 58:6d:8f:04:8a:d3 -> ff:ff:ff:ff:ff:ff ARP 60 Who has 192.168.1.132?  Tell 192.168.1.1

0000  ff ff ff ff ff ff 58 6d 8f 04 8a d3 08 06 00 01   ......Xm........
0010  08 00 06 04 00 01 58 6d 8f 04 8a d3 c0 a8 01 01   ......Xm........
0020  00 00 00 00 00 00 c0 a8 01 84 00 00 00 00 00 00   ................
0030  00 00 00 00 00 00 00 00 00 10 41 1e ff            ..........A..

  1.849838 08:00:27:17:25:16 -> 58:6d:8f:04:8a:d3 ARP 42 192.168.1.132 is at 08:00:27:17:25:16

0000  58 6d 8f 04 8a d3 08 00 27 17 25 16 08 06 00 01   Xm......'.%.....
0010  08 00 06 04 00 02 08 00 27 17 25 16 c0 a8 01 84   ........'.%.....
0020  58 6d 8f 04 8a d3 c0 a8 01 01                     Xm........

  1.852508  65.54.87.10 -> 192.168.1.132 TCP 66 80 > 50968 [SYN, ACK] Seq=0 Ack=1 Win=65535 Len=0 MS:

0000  08 00 27 17 25 16 58 6d 8f 04 8a d3 08 00 45 20   ..'.%.Xm......E
0010  00 34 5d 66 40 00 36 06 8c d1 41 36 57 0a c0 a8   .4]f@.6...A6W...
0020  01 84 00 50 c7 18 9c 36 4c e6 2b 07 39 c7 80 12   ...P...6L.+.9...
0030  ff ff 00 45 00 00 02 04 05 b4 01 03 03 04 04 02   ...E............
0040  00 00                                             ..

  1.852745 192.168.1.132 -> 65.54.87.10  TCP 54 50968 > 80 [ACK] Seq=1 Ack=1 Win=65700 Len=0

0000  58 6d 8f 04 8a d3 08 00 27 17 25 16 08 00 45 00   Xm......'.%...E.
0010  00 28 25 72 40 00 80 06 00 00 c0 a8 01 84 41 36   .(%r@.........A6
0020  57 0a c7 18 00 50 2b 07 39 c7 9c 36 4c e7 50 10   W....P+.9..6L.P.
0030  40 29 5a 87 00 00                                 @)Z...
```

Source: Wireshark

Figure 3-10 A normal exchange of packets between a Web browser and a Web server

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The packets in Figure 3-10 can be broken down as follows:

1. The first packet has the SYN flag set as the Web browser at 192.168.1.132:50969 asks to synchronize a session with the Web server 65.54.87.10:80.
2. Because the Web browser has an IP address in the private address range, it must use a proxy server that can perform NAT (Network Address Translation) to send and receive communications with systems on the Internet, like the Web server it is trying to reach. In the second packet, the proxy server at 192.168.1.1 is broadcasting an ARP query to determine the MAC address of the Web browser's system. Note that the MAC broadcast is FF-FF-FF-FF-FF-FF.
3. In the third packet, the Web browser's system replies to the ARP query with information about its MAC address.
4. Now that the proxy server knows where to find the Web browser, it forwards the second step of the TCP three-way handshake: the SYN ACK from the Web server to the Web browser.
5. The fifth packet shows the Web browser's ACK, which completes the handshake; the Web session is now ready to begin.

Once the handshake is completed, the Web browser sends a request to the Web server for Web page data. This request is called an HTTP GET packet, as shown in Figure 3-11. TShark has labeled this packet as using the HTTP protocol, but remember that the packet must also have a Data Link layer header with source and destination MAC addresses, a Network layer header with source and destination IP addresses, and a Transport layer header with TCP source and destination ports. The figure indicates that the GET packet contains 399 bytes, that the Web browser supports HTTP version 1.1, and that the Web site being accessed is *www.msn.com*.

```
 1.855278 192.168.1.132 -> 65.54.87.10  HTTP 399 GET /i/52/79D26BA7D3D2DF25DBDD991B35E16.jpg HTTP/1.1

0000  58 6d 8f 04 8a d3 08 00 27 17 25 16 08 00 45 00   Xm......'.%...E.
0010  01 81 25 75 40 00 80 06 00 00 c0 a8 01 84 41 36   ..%u@.........A6
0020  57 0a c7 17 00 50 a7 1c 46 4f d3 06 5e 67 50 18   W....P..FO..^gP.
0030  40 29 5b e0 00 00 47 45 54 20 2f 69 2f 35 32 2f   @)[...GET /i/52/
0040  37 39 44 32 36 42 41 37 44 33 44 32 44 46 32 35   79D26BA7D3D2DF25
0050  44 42 44 44 39 39 31 42 33 35 45 31 36 2e 6a 70   DBDD991B35E16.jp
0060  67 20 48 54 54 50 2f 31 2e 31 0d 0a 41 63 63 65   g HTTP/1.1..Acce
0070  70 74 3a 20 69 6d 61 67 65 2f 70 6e 67 2c 20 69   pt: image/png, i
0080  6d 61 67 65 2f 73 76 67 2b 78 6d 6c 2c 20 69 6d   mage/svg+xml, im
0090  61 67 65 2f 2a 3b 71 3d 30 2e 38 2c 20 2a 2f 2a   age/*;q=0.8, */*
00a0  3b 71 3d 30 2e 35 0d 0a 52 65 66 65 72 65 72 3a   ;q=0.5..Referer:
00b0  20 68 74 74 70 3a 2f 2f 77 77 77 2e 6d 73 6e 2e    http://www.msn.
00c0  63 6f 6d 2f 3f 6f 63 69 64 3d 69 65 68 70 0d 0a   com/?ocid=iehp..
00d0  41 63 63 65 70 74 2d 4c 61 6e 67 75 61 67 65 3a   Accept-Language:
00e0  20 65 6e 2d 55 53 0d 0a 55 73 65 72 2d 41 67 65    en-US..User-Age
00f0  6e 74 3a 20 4d 6f 7a 69 6c 6c 61 2f 35 2e 30 20   nt: Mozilla/5.0
0100  28 63 6f 6d 70 61 74 69 62 6c 65 3b 20 4d 53 49   (compatible; MSI
0110  45 20 39 2e 30 3b 20 57 69 6e 64 6f 77 73 20 4e   E 9.0; Windows N
0120  54 20 36 2e 31 3b 20 57 4f 57 36 34 3b 20 54 72   T 6.1; WOW64; Tr
0130  69 64 65 6e 74 2f 35 2e 30 29 0d 0a 41 63 63 65   ident/5.0)..Acce
0140  70 74 2d 45 6e 63 6f 64 69 6e 67 3a 20 67 7a 69   pt-Encoding: gzi
0150  70 2c 20 64 65 66 6c 61 74 65 0d 0a 48 6f 73 74   p, deflate..Host
0160  3a 20 63 6f 6c 2e 73 74 62 2e 73 2d 6d 73 6e 2e   : col.stb.s-msn.
0170  63 6f 6d 0d 0a 43 6f 6e 6e 65 63 74 69 6f 6e 3a   com..Connection:
0180  20 4b 65 65 70 2d 41 6c 69 76 65 0d 0a 0d 0a       Keep-Alive....
```

Source: Wireshark

Figure 3-11 An HTTP GET packet

Examining Abnormal Network Traffic Signatures

As IDPSs become more sophisticated, the techniques that attackers use to circumvent them have multiplied and become more complicated. Features such as illegal combinations of TCP flags and private IP addresses in packets are relatively easy to identify as abnormal,

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

compared with attacks that use a range of packets. Suspicious traffic signatures can fall into one of the following categories:

- *Informational*—This traffic might not be malicious, but it could be used to verify whether an attack has been successful. Examples include ICMP echo request packets or TCP packets sent to a specific port on a specific system.
- *Reconnaissance*—This traffic could represent an attacker's attempt to gain information about a network as a prelude to an attack. Examples include ping sweeps and port scans.
- *Unauthorized access*—This traffic might be caused by someone who has gained unauthorized access to a system and is attempting to retrieve data from it. Examples include the Shixploit attack and the Hydraq attack.
- *Denial of service*—This traffic might be part of an attempt to slow or halt all connections on a network device, such as a Web server or mail server. ARP cache poisoning is an example of an attack that can cause denial of service.

Some of the more common examples of suspicious traffic—ping sweeps, port scans, random back door scans, and Trojan scans—are described in the following sections along with their signatures.

Ping Sweeps To gain access to specific resources on an internal network, a hacker needs to determine the location of a host. One method is to conduct a **ping sweep** (also called an ICMP sweep), which sends a series of ICMP echo request packets in a range of IP addresses. Usually, the messages come in quick succession. Multiple packets can be detected in a single second, indicating that an automated tool is being used. An example of a ping sweep is shown in Figure 3-12. Notice that the source is always the same system but the destination IP addresses change sequentially. Be sure to examine the PING times in the first column; the sweep illustrated in the figure took place in about two seconds.

```
2.717757 192.168.1.110 -> 67.9.155.1   ICMP 42 Echo (ping) request  id=0x9188, seq=0/0, ttl=39
2.717914 192.168.1.110 -> 67.9.155.2   ICMP 42 Echo (ping) request  id=0xab94, seq=0/0, ttl=42
2.718227 192.168.1.110 -> 67.9.155.3   ICMP 42 Echo (ping) request  id=0x0ef3, seq=0/0, ttl=45
2.718309 192.168.1.110 -> 67.9.155.4   ICMP 42 Echo (ping) request  id=0x08e6, seq=0/0, ttl=50
2.718402 192.168.1.110 -> 67.9.155.5   ICMP 42 Echo (ping) request  id=0x50c8, seq=0/0, ttl=53
2.718487 192.168.1.110 -> 67.9.155.6   ICMP 42 Echo (ping) request  id=0xddc7, seq=0/0, ttl=48
2.718580 192.168.1.110 -> 67.9.155.7   ICMP 42 Echo (ping) request  id=0xb78a, seq=0/0, ttl=38
2.718654 192.168.1.110 -> 67.9.155.8   ICMP 42 Echo (ping) request  id=0xbcd2, seq=0/0, ttl=50
2.718723 192.168.1.110 -> 67.9.155.9   ICMP 42 Echo (ping) request  id=0xb453, seq=0/0, ttl=52
2.718792 192.168.1.110 -> 67.9.155.10  ICMP 42 Echo (ping) request  id=0xd793, seq=0/0, ttl=52
2.719075 192.168.1.110 -> 67.9.155.11  ICMP 42 Echo (ping) request  id=0xdc11, seq=0/0, ttl=39
2.719166 192.168.1.110 -> 67.9.155.12  ICMP 42 Echo (ping) request  id=0xf91c, seq=0/0, ttl=45
3.730437 192.168.1.110 -> 67.9.155.13  ICMP 42 Echo (ping) request  id=0xb7b9, seq=0/0, ttl=52
3.730716 192.168.1.110 -> 67.9.155.14  ICMP 42 Echo (ping) request  id=0x9f11, seq=0/0, ttl=53
3.730919 192.168.1.110 -> 67.9.155.15  ICMP 42 Echo (ping) request  id=0x65df, seq=0/0, ttl=58
3.731091 192.168.1.110 -> 67.9.155.16  ICMP 42 Echo (ping) request  id=0x6cf1, seq=0/0, ttl=43
3.731260 192.168.1.110 -> 67.9.155.17  ICMP 42 Echo (ping) request  id=0x2230, seq=0/0, ttl=52
3.731672 192.168.1.110 -> 67.9.155.18  ICMP 42 Echo (ping) request  id=0x5b79, seq=0/0, ttl=48
3.732010 192.168.1.110 -> 67.9.155.19  ICMP 42 Echo (ping) request  id=0xea59, seq=0/0, ttl=45
3.732192 192.168.1.110 -> 67.9.155.20  ICMP 42 Echo (ping) request  id=0xe1e6, seq=0/0, ttl=39
3.732437 192.168.1.110 -> 67.9.155.21  ICMP 42 Echo (ping) request  id=0xab55, seq=0/0, ttl=44
3.732800 192.168.1.110 -> 67.9.155.22  ICMP 42 Echo (ping) request  id=0xbd48, seq=0/0, ttl=47
```

Source: Wireshark

Figure 3-12 An automated ping sweep

The ping sweep, by itself, does not cause harm to computers on your network. The prudent response is to note the IP address used in the ping sweep to track further activity. An IDPS could be configured to transmit an alarm and block transmissions if the IP address in question attempts to connect to a specific host on the network, for instance.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Port Scans If an attacker can determine any legitimate IP addresses on an internal network, the next step is to target one of those addresses and perform a **port scan**—an attempt to connect to a computer's ports and check whether any are active and listening. An attacker who finds an open port can exploit any known vulnerabilities associated with any service that runs on that port. Many newer applications perform dynamic port negotiation based on available resources. However, vulnerabilities still exist in older applications and services.

The signature of a port scan typically includes a SYN packet sent in sequence to each port on an IP address, as shown in Figure 3-13.

```
0.048185 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 1  [SYN] Seq=0 Win=3072 Len=0 MSS=1460
0.048323 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 2  [SYN] Seq=0 Win=4096 Len=0 MSS=1460
0.048405 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 3  [SYN] Seq=0 Win=3072 Len=0 MSS=1460
0.048484 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 4  [SYN] Seq=0 Win=3072 Len=0 MSS=1460
0.048563 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 5  [SYN] Seq=0 Win=1024 Len=0 MSS=1460
0.048643 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 6  [SYN] Seq=0 Win=3072 Len=0 MSS=1460
0.048723 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 7  [SYN] Seq=0 Win=3072 Len=0 MSS=1460
0.048801 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 8  [SYN] Seq=0 Win=4096 Len=0 MSS=1460
0.048882 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 9  [SYN] Seq=0 Win=2048 Len=0 MSS=1460
0.049164 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 10 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
0.049593 192.168.1.110 -> 192.168.1.132 TCP 54 45571 > 11 [SYN] Seq=1 Win=1024 Len=0 MSS=1460
0.054744 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 12 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
0.054841 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 13 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
0.054923 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 14 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
0.055003 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 15 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
0.056392 192.168.1.110 -> 192.168.1.132 TCP 54 45571 > 16 [SYN] Seq=0 Win=2048 Len=0 MSS=1460
0.084972 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 17 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
0.085140 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 18 [SYN] Seq=0 Win=4096 Len=0 MSS=1460
0.085241 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 19 [SYN] Seq=0 Win=3072 Len=0 MSS=1460
0.085340 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 20 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
0.085437 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 21 [SYN] Seq=0 Win=2048 Len=0 MSS=1460
0.085539 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 22 [SYN] Seq=0 Win=2048 Len=0 MSS=1460
0.085636 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 23 [SYN] Seq=0 Win=3072 Len=0 MSS=1460
0.085734 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 24 [SYN] Seq=0 Win=3072 Len=0 MSS=1460
0.100816 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 25 [SYN] Seq=0 Win=2048 Len=0 MSS=1460
0.100947 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 26 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
0.101018 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 27 [SYN] Seq=0 Win=2048 Len=0 MSS=1460
0.101086 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 21 [SYN] Seq=0 Win=2048 Len=0 MSS=1460
0.101154 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 28 [SYN] Seq=0 Win=4096 Len=0 MSS=1460
0.101222 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 29 [SYN] Seq=0 Win=2048 Len=0 MSS=1460
0.101290 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 30 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
0.101357 192.168.1.110 -> 192.168.1.132 TCP 58 45571 > 31 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
```

Source: Wireshark

Figure 3-13 An automated port scan

In Figure 3-13, the source port (45571) does not change from packet to packet. This indicates that the attacker is not expecting the target computer to establish a full three-way connection, but wants to find out whether the IP address is being used. In another type of port scan, the source port changes with each packet along with the destination port because a full three-way connection is expected. Notice also that the target IP address stays the same in each packet; a single computer is the target, unlike the ping sweep described earlier in this section. However, the target port is incremented by one on each succeeding packet. Some automated ping sweeps make themselves less obvious by jumping from port to port rather than operating in sequence. The scan shown in the figure is not particularly stealthy because the events are happening so rapidly; the 31 ports were scanned in less than a second.

Random Back Door Scans You can think of a port as a virtual door through which data can enter and leave a computer. In that context, a **back door** is an undocumented or unauthorized hidden opening (such as a port) through which an attacker can access a computer, program, or other resource. One type of port scan probes the same computer ports used by well-known Trojan programs to see if any ports are open and listening. These applications seem to be harmless, but they can damage a computer or the files on it.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Specific Trojan Scans Attackers can execute port scans in several ways. In a **vanilla scan**, all ports from 0 to 65,535 are probed one after another. In a **strobe scan**, an attacker scans only ports that are commonly used by specific programs in an attempt to see whether the program is present and can be used. Table 3-2 shows some examples of Trojans and the protocols and ports associated with them.

Trojan	Protocol	Port
Trojan.Asprox	TCP	80, 82
W32.Spybot.pen	TCP	8076
w32.mytob.jw@mm	TCP	10027
Trojan.Mitglieder.h	TCP	14247
Sub-7 2.1	TCP	27573
Remote Windows Shutdown	TCP	53001
Back Orifice 2000	UDP	54321

© Cengage Learning 2014

Table 3-2 Examples of Trojan programs and ports

A common type of strobe scan searches IP addresses on a network for the presence of a specific Trojan program. If attackers can find a Trojan program that has already circumvented the firewall and IDPS and is already operating, they can save themselves the time and effort of installing a new Trojan program. For instance, in Figure 3-14, a target host at 192.168.1.132 is being scanned on several ports that are commonly used by known Trojans. An analysis of this capture follows:

- *Packet 1*—The scanner at 192.168.1.110 sends a SYN packet to the target's port 4646.
- *Packet 2*—The scanner sends a second SYN packet to the target, this time to port 5401.
- *Packet 3*—The target responds to the first packet with an RST ACK packet.
- *Packet 4*—The scanner sends a SYN to the target's port 4512.
- *Packets 5 and 6*—The target responds to the scanner's second and third probes with RST ACK packets.

Note that the scanner is always identifiable by its IP address, but it randomly changes the source port from which it sends its SYN probes.

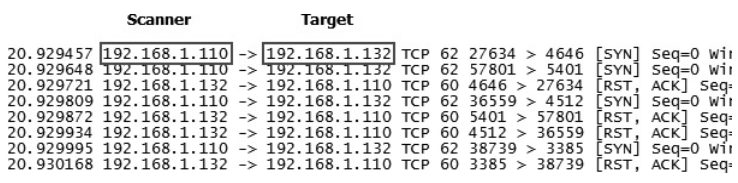

Source: Wireshark

Figure 3-14 A scan of a single host for existing Trojans

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

If you see this type of scan in your log files, you need to act quickly to block the source IP address because attacks probably will occur in the near future. You will learn about blocking specific IP addresses, subnets, or address ranges in Chapters 9 and 10 when you study firewalls.

Many newer Trojans do not listen at a predetermined port on the victim's computer, but generate a random port address for listening. This approach makes detection more difficult, but the attacker also encounters more trouble finding a back door that is already installed.

Antivirus software and many intrusion-detection systems have something in common: They both look for known virus and attack signatures. They also both require up-to-date signature files to operate at peak efficiency.

Nmap Scans Network Mapper (Nmap) is a popular software tool for scanning networks. You should be able to recognize the common types of scans it enables attackers to perform. With Nmap, attackers can send packets that circumvent the normal three-way handshakes two computers use to establish a connection. Nmap enables attackers to send packets for which an IDPS might not be configured to send an alarm. The IDPS might see a combination of TCP flags that it does not recognize, and because no rule exists for the combination, an alarm might not be triggered. Examples of Nmap scans include the following:

- *SYN scans*—The attacker sends a progression of packets with only the SYN flag set. The targeted computer responds with packets that have the ACK flag set, but the originating computer simply keeps sending SYN packets. See Figure 3-15.

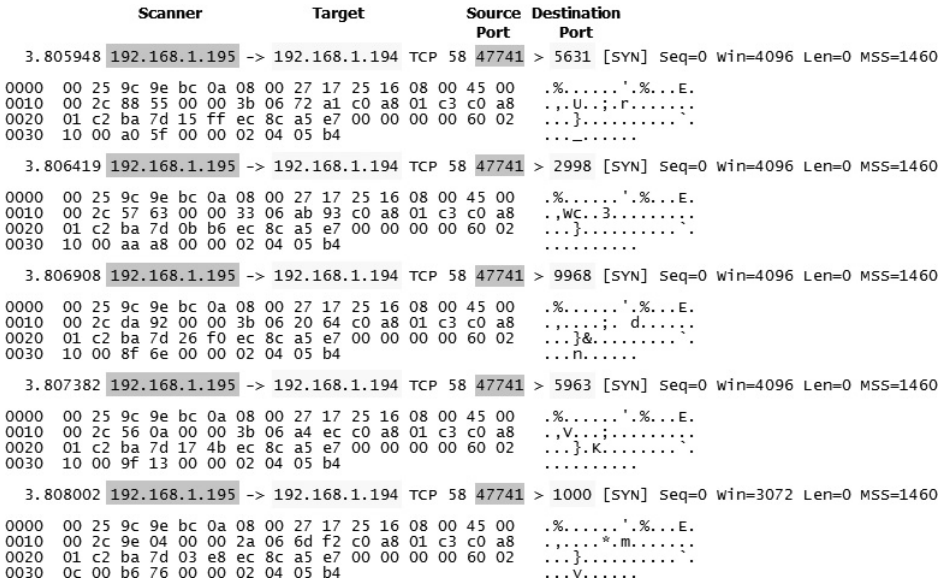

Source: Nmap

Figure 3-15 Nmap SYN scan

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *FIN scans*—The attacker sends only packets that have the FIN flag set; a SYN flag is never sent.
- *ACK scans*—The attacker sends only packets with the ACK flag set; a SYN or FIN flag is never sent.
- *Null scans*—The attacker sends a sequence of packets that have no flags set. An IDPS might ignore packets with no flags set.
- *Xmas scans*—The attacker sends a sequence of packets that have the FIN PSH URG flags set. In Figure 3-16, the scanner at 192.168.1.195 is scanning both 192.168.1.1 and 192.168.1.194 from the same source port, but is "randomizing" the target ports. However, in the first and fourth packets, you can see that the scanner has used the same target port on the two target computers.

```
  6.420376 192.168.1.195 -> 192.168.1.1   TCP 54 35048 > 6389 [FIN, PSH, URG] Seq=1 Win=1024 Urg=0 Len=0

0000  30 46 9a 15 33 bb 08 00 27 17 25 16 08 00 45 00   0F..3...'.%...E.
0010  00 28 f2 46 00 00 30 06 14 75 c0 a8 01 c3 c0 a8   .(.F..0..u......
0020  01 01 88 e8 18 f5 af 53 4d 2d 00 00 00 00 50 29   .......SM-....P)
0030  04 00 89 48 00 00                                 ...H..

  6.420999 192.168.1.195 -> 192.168.1.194 TCP 54 35048 > 8651 [FIN, PSH, URG] Seq=1 Win=2048 Urg=0 Len=0

0000  00 25 9c 9e bc 0a 08 00 27 17 25 16 08 00 45 00   .%......'.%...E.
0010  00 28 44 01 00 00 2d 06 c4 f9 c0 a8 01 c3 c0 a8   .(D...-.........
0020  01 c2 88 e8 21 cb af 53 4d 2d 00 00 00 00 50 29   ....!..SM-....P)
0030  08 00 7b b1 00 00                                 ..{...

  6.421224 192.168.1.195 -> 192.168.1.1   TCP 54 35048 > 2004 [FIN, PSH, URG] Seq=1 Win=4096 Urg=0 Len=0

0000  30 46 9a 15 33 bb 08 00 27 17 25 16 08 00 45 00   0F..3...'.%...E.
0010  00 28 90 98 00 00 37 06 6f 23 c0 a8 01 c3 c0 a8   .(....7.o#......
0020  01 01 88 e8 07 d4 af 53 4d 2d 00 00 00 00 50 29   .......SM-....P)
0030  10 00 8e 69 00 00                                 ...i..

  6.421508 192.168.1.195 -> 192.168.1.194 TCP 54 35048 > 6389 [FIN, PSH, URG] Seq=1 Win=3072 Urg=0 Len=0

0000  00 25 9c 9e bc 0a 08 00 27 17 25 16 08 00 45 00   .%......'.%...E.
0010  00 28 34 b5 00 00 2a 06 d7 45 c0 a8 01 c3 c0 a8   .(4...*..E......
0020  01 c2 88 e8 18 f5 af 53 4d 2d 00 00 00 00 50 29   .......SM-....P)
0030  0c 00 80 87 00 00                                 ......

  6.421803 192.168.1.195 -> 192.168.1.1   TCP 54 35048 > 1054 [FIN, PSH, URG] Seq=1 Win=3072 Urg=0 Len=0

0000  30 46 9a 15 33 bb 08 00 27 17 25 16 08 00 45 00   0F..3...'.%...E.
0010  00 28 5d 2e 00 00 3a 06 9f 8d c0 a8 01 c3 c0 a8   .(]...:.........
0020  01 01 88 e8 04 1e af 53 4d 2d 00 00 00 00 50 29   .......SM-....P)
0030  0c 00 96 1f 00 00                                 ......
```

Source: Nmap

Figure 3-16 Nmap Xmas scan

In each of these scans, a three-way handshake can never be established with the computer attempting to make a connection. The attacker probably is attempting to determine whether an application is active on a particular port.

Identifying Suspicious Events

The previous section described well-known attack signatures you might encounter when inspecting your IDPS log files. Unfortunately, attackers often avoid launching such well-known attacks. Instead, they use more subtle means to try to gain unauthorized access to computers in your network. For instance, instead of the sequence of FIN packets described in the previous section, you might see only a single FIN packet sent to a port on a computer. Other "orphaned" packets might follow, but only after an interval during which a substantial

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

number of legitimate packets have passed through. The use of an interval acts as a distraction during which the different stages of the attack are separated, so an administrator might not realize that they are part of the same attack. **Scan throttling** is often used by attackers to delay the progression of a scan over hours, days, or even weeks. For example, the scans you perform in the Hands-On Projects at the end of this chapter should take less than a minute, but they might take a month when controlled by a cautious attacker of a targeted network. You can configure newer IDPSs to track this kind of activity, and you can use variable dynamic throttling so that the security device can adjust the threshold that triggers a warning or action, depending on the rate of detected scans.

These attacks can be extremely difficult to detect if you review log files manually as part of your job. Log files can quickly amass so many entries that they become overwhelming to review, and you miss important entries. Fortunately, you can avoid the tedious chore of poring through thousands of log entries for known attack signatures or other abnormal signatures by using an extensive database of signatures that includes these events.

This section describes events and characteristics of network communications that you need to identify as suspicious after your IDPS has responded to them by transmitting an alarm. The discussion includes packet header discrepancies, advanced IDPS attacks, and Remote Procedure Calls (RPC) abuses.

Packet Header Discrepancies

Discrepancies you find in TCP, IP, ICMP, or UDP packet headers can provide warning signs that an attacker has crafted the packet (in other words, manufactured or altered it on purpose). However, instead of seeing these discrepancies in a well-defined and lengthy succession of packets, you might receive only a single packet with a falsified IP address, falsified port number, illegal TCP flags, illegal TCP or IP options, or fragmentation abuses.

Falsified IP Address Your IDPS might send alarms for violations of IP header settings, as specified in RFC 791, "Internet Protocol." For example, an IP address should not appear in one of the three reserved ranges (10.0.0.0 to 10.255.255.255, 172.16.0.0 to 172.31.255.255, and 192.168.0.0 to 192.168.255.255). The use of addresses in the reserved ranges is limited to private networks. If you see one of these addresses in a packet coming from the public network, a router or other device might have been misconfigured or is malfunctioning. On the other hand, the packet might show a private address because an attacker has used **IP spoofing**. In other words, the attacker has inserted a false address into the IP header to make the packet more difficult to trace back to its source.

For more information about the correct use of private IP addresses, refer to RFC 1918, "Address Allocation for Private Internets," at *www.faqs.org/rfcs/rfc1918.html*.

A land attack is an example of a falsified IP address used to cause a server to malfunction. It occurs when a detected IP packet has an invalid IP address setting in which the source and destination IP/port addresses are the same—the target's IP address. Another attack that uses an invalid IP address, the localhost source spoof, should trigger an alarm if the local host source address of 127.0.0.1 occurs in a packet.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Falsified Port Number or Protocol You already know that IP address information can be falsified in a packet, but protocol numbers might also be altered to elude an IDPS. TCP and UDP headers should never have the source or destination port set to 0 because this number is reserved by IANA. Protocol numbers are assigned by IANA, like port numbers. The Protocol field in IPv4 is an 8-bit field used to specify this number. (In IPv6, the field is called Next Header.) This field indicates the transport protocol to be used, such as ICMP (#1), TCP (#6), or UDP (#17). Currently, the number in this field cannot be set greater than 142. The numbers 143–252 are unassigned, 253 and 254 are used for experimentation and testing, and 255 is reserved. The use of undefined protocol numbers might indicate an attacker's attempt to establish a proprietary communications channel—a channel known to one person and used only by that person.

Illegal TCP Flags As you saw in the section called "Examining Normal Network Traffic Signatures" earlier in this chapter, the TCP flags SYN and ACK are exchanged to establish a connection between two computers. The PSH flag is used when data is being sent, and the FIN flag is used when a connection is complete. Other normal TCP flag rules include the following:

- Every packet in a connection should have the ACK bit set, except for the initial SYN packet and possibly an RST packet used to terminate a connection.
- Packets during the "conversation" portion of the connection contain just an ACK flag by default. This portion of the connection occurs after the three-way handshake but before teardown or termination. Optionally, these packets can contain PSH and/or URG flags.
- FIN/ACK and ACK are used during the normal teardown of an existing connection. PSH, FIN, and ACK might also be seen near the end of a connection.
- RST or RST/ACK can be used to terminate a connection immediately.

One of the most obvious ways to detect an abnormal packet signature is to look at the TCP flags for violations of normal usage. A packet with the SYN and FIN flags set should not exist in normal traffic; however, an attacker might set both flags to cause the destination computer to crash or freeze because it does not know how to respond. After the server is disabled, the attacker can attack a computer on the internal network using an IP address detected earlier through network scans.

The following list summarizes signatures of malformed packets that misuse the SYN and FIN flags:

- SYN FIN is probably the best-known illegal combination. Because SYN is used to start a connection and FIN is used to end one, it does not make sense to include both flags together in a packet. Many scanning tools use SYN FIN packets because older intrusion-detection systems often were not configured to recognize or block them. However, most IDPS devices are now configured to catch such illegal combinations. You can safely assume that hackers created any SYN FIN packets you see.
- Other variants of SYN FIN exist, including SYN FIN PSH, SYN FIN RST, and SYN FIN RST PSH. Their use is sometimes called an Xmas attack. These packets can be used by attackers who know that IDPSs might be looking for packets with just the SYN and FIN flags set.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Packets should never contain a FIN flag by itself. FIN packets are frequently used for port scans, network mapping, and other stealth activities.
- A SYN-only packet, which should occur only when a new connection is being initiated, should not contain any data.

You might also encounter **null packets**—TCP packets with no flags set, which could cause a server to crash. It is a violation of TCP rules to use a packet with no flags set.

TCP or IP Options TCP options in a packet can alert you to intrusion attempts and even enable you to identify the type of OS being used. For instance, only one MSS or window option should appear in a packet. MSS, NOP, and SackOK should appear only in packets that have the SYN and/or ACK flag set.

IPv4 options were originally intended as ways to insert special handling instructions into packets. However, attackers mostly use IP options now for attack attempts. Because of this vulnerability, many filters simply drop all packets with IPv4 options set. In IPv6, the options field is removed and replaced by extension headers.

Fragmentation Abuses Every type of computer network, such as Ethernet, FDDI, and Token Ring, has its own **maximum transmission unit** (**MTU**)—the maximum packet size it can transmit. Packets that are larger than the MTU must be fragmented, or broken into multiple segments that are small enough for the network to handle.

After a packet is broken into fragments, each fragment receives its own IP header. However, in IPv4, only the initial packet in a set includes a header for higher-level protocols. Most filters need the information in the higher-level protocol header to make the decision to allow or deny the packet. Accordingly, attackers send only secondary fragments, which are any fragments other than the initial one. These packets are often allowed past the IDPS because filter rules are applied to first fragments only. IPv6 addresses this vulnerability in two ways: It permits only the source node to fragment payloads, and it divides the packet into unfragmentable and fragmentable parts. The fragmentable part can be processed only when the packet reaches its destination. The unfragmentable portion includes the IPv6 header, the hop-by-hop options header, the destination options header, and the routing header. This portion is processed by each router along the route to the destination.

Fragmentation can occur normally. However, an IDPS should be configured to send an alarm if it encounters a large number of fragmented packets. Many different types of fragmentation abuses can occur. Some of the more serious abuses are described briefly in the following list:

IPv4

- *Overlapping fragments*—Two fragments of the same packet have the same position within the packet, so the contents overlap. A properly configured firewall should always drop this type of packet.
- *Fragments that are too large*—An IP packet can be no larger than 65,535 bytes. If packets are more than the maximum size when they are reassembled from their fragments, they might cause some systems to crash. This activity could indicate a DoS attack.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *Fragments overwrite data*—Some early fragments in a sequence are transmitted along with random data. Later fragments overwrite the random data. If the packet is not reassembled properly, the IDPS cannot detect the attack.
- *Fragments are too small*—If any fragment other than the final fragment in a sequence is less than 400 bytes, it has probably been crafted intentionally. Such a small fragment is probably part of a DoS attack.

IPv6

- *Fragments with a destination address of a network device*—Assembly of fragmented packets occurs only at the destination host. If a router, firewall, or other device is the destination of fragmented IPv6 packets, a denial of service attack might be intended.
- *Fragments are too small*—If any fragment other than the final fragment in a sequence is less than 1280 bytes, it has probably been crafted intentionally.
- *Fragments that arrive too slowly*—Fragments that take more than 60 seconds to deliver should be dropped because they are probably part of an effort to avoid detection.

Advanced Attacks

Most types of attacks discussed so far have been protocol anomalies—violations of the protocol rules described in RFC statements. Some especially complex attacks use path names, hexadecimal codes, and obfuscated directory names to fool an IDPS into letting the packet through without triggering an alarm. Some of the more advanced IDPS evasion techniques include the following:

- *Polymorphic buffer overflow attacks*—These attacks are as complicated as they sound. A tool called ADMmutate is used to alter an attack's shell code so that it differs slightly from the known signatures many IDPSs use. After the attacking packets elude the IDPS and reach their intended target, they reassemble into their original form.
- *Path obfuscation*—A directory path statement in the payload of a packet is obfuscated by using multiple forward slashes. For example, */Windows/. /. /. /* is essentially the same as */Windows*. However, because the signatures do not match exactly, an IDPS might be unable to detect this attack.
- *CGI scripts*—A series of packets is sent to a series of well-known **Common Gateway Interface (CGI) scripts**, which are scripts used to process data submitted over the Internet. Examples include CGI scripts such as Count.cgi, FormMail, AnyForm, Php.cgi, TextCounter, and GuestBook. You can be certain that someone is attempting to exploit your network if it does not contain these files, but packets attempt to locate them anyway.
- ***Packet injection***—With readily available tools like Nemesis, attackers can easily craft packets that comply with protocols like ARM, DNS, ICMP, Ethernet, IGMP, IP, OSPF, RIP, TCP, and UDP. These packets can be inserted into network traffic. While these tools are useful for testing IDPSs and firewalls, they can also be used to disrupt communications, spoof a variety of systems, and carry out a number of attacks.

The only way to avoid these attacks is to keep your IDPS signatures up to date and to watch your log files closely.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Remote Procedure Call Attacks

Remote Procedure Calls (**RPC**) is a standard set of communications rules that allows one computer to request a service (in other words, a remote procedure) from another computer on a network. RPC uses a program called a portmapper that maintains a record of each remotely accessible program and the port it uses. The RPC portmapper is actually a service that runs on the system and converts RPC program numbers into TCP/IP protocol port numbers. Because RPC can provide remote access to applications, attackers naturally attempt to use it to gain unauthorized access to those applications. Here are some examples of RPC-related events that should trigger IDPS alarms:

- *RPC dump*—A targeted host receives an RPC dump request: a request to report the presence and port use of any RPC services the system provides.
- *RPC set spoof*—A targeted host receives an RPC set request from a source IP address of 127.0.0.1.
- *RPC NFS sweep*—A targeted host receives a series of requests for the NFS program on a succession of different ports.

RPC services such as Network Information System (NIS) use a four-byte service number because there are too many services to use a two-byte port number. When an RPC service starts, it allocates a random TCP or UDP port for itself. It then contacts rpcbind or portmapper and registers its service number and TCP/UDP port. Portmapper and rpcbind always run on port 111, for example. A client that wants to talk to a server first contacts portmapper to get the port number, and then continues the exchange with the server directly. A client can bypass portmapper and scan for services. There is no guarantee that a particular service will end up on a particular port.

The rpcbind daemon is a more recent implementation of portmapper. It includes all of the functionality of portmapper and additional features.

Chapter Summary

- Network security hardware and software should work cooperatively to share information. A standard called Common Vulnerabilities and Exposures (CVE) enables IDPSs, firewalls, and other devices to share attack signatures and information about network vulnerabilities so that they can better protect networks. Mitre maintains a list of current vulnerabilities as an online database; you can use the list to update your own CVE database and learn about new attacks.
- Interpreting the signatures of normal and abnormal network traffic can help prevent network intrusions. Recognizing the characteristics of a possible intrusion makes it possible to interpret and react effectively to log files and alert messages.
- You can adjust filter rules to reduce the number of false alarms you receive from your IDPS. More importantly, you can prevent intrusions before they occur or keep intrusions that are already under way from causing excessive damage.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- The analysis of traffic signatures is an integral aspect of intrusion prevention. A signature is a set of characteristics, such as IP addresses, port numbers, TCP flags, and options, that define network activity. Normal traffic makes valid use of these settings. Possible intrusions are marked by invalid settings.
- Suspicious TCP/IP packets can include bad header information, suspicious contents in the data payload of packets, IP options settings, and a succession of packets, such as a denial of service attempt.
- TCP flags are used in sequence to create a normal three-way handshake between two computers.
- By learning what normal traffic signatures look like, you can identify signatures of suspicious connection attempts. You can monitor suspicious events, such as ping sweeps, port scans, random back door scans, and Trojan scans. You should also be familiar with the characteristics of packets crafted with Nmap.
- You can identify a variety of other suspicious network events, including "orphaned" packets, land attacks in which the source and destination IP/port addresses are the same, localhost source spoofs, falsified protocol numbers, and illegal combinations of TCP flags, such as SYN/FIN.
- Advanced attacks are especially difficult to detect without a database of intrusion signatures or user behaviors. Certain complex attacks called polymorphic buffer overflow attacks can be altered so that they do not match a known intrusion signature and elude the IDPS. Other complex attacks use confusing path names or other keywords in the data payload section of packets.
- Other advanced attack methods might include attempts to abuse common CGI scripts or misuse Remote Procedure Calls that enable remote users to access services on a computer.

Key Terms

back door An undocumented or unauthorized hidden opening that permits access to a system.

checksum A simple error-checking procedure used to determine whether a message has been tampered with in transit.

Common Gateway Interface (CGI) scripts Scripts used to process data submitted over the Internet.

Common Vulnerabilities and Exposures (CVE) A standard that enables security devices to share information about attack signatures and other vulnerabilities so that they can work together to protect networks.

ICMP ID number A number that identifies the ICMP packet so that the originating computer can make sure the response came from its original request.

IP spoofing The process of inserting a false address into an IP header to make the packet more difficult to trace back to its source.

maximum transmission unit (MTU) The maximum packet size that can be transmitted over a type of computer network. For example, the MTU of an Ethernet network is 1500 bytes.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

multiple-packet attacks Attacks that require a series of packets to be transmitted in order to execute.

null packets TCP packets with no flags set.

packet injection The insertion of crafted packets into network traffic.

packet sniffer Software or hardware that monitors traffic into or out of network devices and captures information about each TCP/IP packet it detects.

ping sweep The act of sending a series of ICMP echo request packets in a range of IP addresses to see whether any computers respond.

port scan An attempt to connect to a computer's ports to see if any are active and listening.

remote-access Trojans (RATs) Malicious programs that allow a remote attacker to take control of the target system.

Remote Procedure Calls (RPC) A standard set of communication rules that allows one computer to request a service from another computer on a network.

scan throttling The tactic of extending the time between scan probes to make the process more difficult to detect.

scanner A device that scans a network for open ports or other potential vulnerabilities.

selective acknowledgements Acknowledgements that selected packets in a sequence have been received, as opposed to acknowledging every packet.

signature A set of characteristics—such as IP numbers and options, TCP flags, and port numbers—used to define a type of network activity.

signature analysis The practice of analyzing and understanding TCP/IP communications to determine whether traffic is legitimate or suspicious.

single-packet attack An attack that can be completed by sending a single network packet from client to host.

strobe scan A type of port scan that probes ports commonly used by specific programs and attempts to use the programs. A common type of strobe scan looks for active Trojan programs that have already circumvented security measures and are running on the scanned system. If an attacker can find an installed Trojan, he can use it instead of having to install a new one.

vanilla scan A type of port scan in which all ports from 0 to 65,535 are probed, one after another.

Review Questions

1. Security devices on a network process digital information, such as text files and Web pages, in the same way. However, which of the following pieces of information might they handle differently?
   a. protocols
   b. TCP/IP headers
   c. attack signatures
   d. port numbers

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

2. In which of the following situations can CVE improve the coordination of intrusion information on a network?
   a. Attack signatures can be compared to lists of known attack signatures on the CVE Web site.
   b. Attack signatures can be compared to current network topology.
   c. Installing application patches can thwart a reported attack.
   d. Current network vulnerabilities can be used to generate application patches.
3. Which of the following can be included in a network traffic signature? (Choose all that apply.)
   a. logon attempts
   b. message digest
   c. TCP options
   d. Ethernet interface number
4. What is the name of an error-checking procedure that uses a formula to calculate a numeric value?
   a. check string
   b. one-way hash
   c. hexadecimal code
   d. checksum
5. How do attackers use fragmentation to circumvent network defenses? (Choose all that apply.)
   a. Fragments are crafted to be too large or too small.
   b. The initial packet is missing.
   c. The fragments arrive too slowly.
   d. The final fragment sent is less than 400 bytes long.
6. Which of the following packets should never have a data payload?
   a. one with SYN/ACK flags set
   b. one with the ACK flag set
   c. one with the SYN flag set
   d. one with ACK/PSH flags set
7. Which of the following is not required for a single-packet attack? (Choose all that apply.)
   a. a source IP address
   b. a destination IP address
   c. an ICMP echo request
   d. an established connection

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

8. Which of the following is an example of a multiple-packet attack?
   a. a fragment
   b. an ICMP flood
   c. a false Internet time stamp
   d. a packet with SYN/FIN/ACK flags set

9. What is the purpose of the 4-byte acknowledgement number in a TCP header?
   a. It acknowledges receipt of the previous packet in the sequence.
   b. It acknowledges that a connection has been made.
   c. It verifies that the source and destination IP addresses are correct.
   d. It acknowledges the ID number the packet is using.

10. Which of the following is the correct order in which TCP flags appear during the initiation of a normal connection?
    a. SYN, ACK, FIN, RST
    b. SYN, PSH, ACK, RST
    c. SYN, SYN/ACK, ACK
    d. SYN, PSH, ACK, FIN

11. Which protocol uses one port number to establish a connection and a different port number to transfer data?
    a. TCP/IP
    b. FTP
    c. HTTP
    d. ICMP

12. Which of the following is an example of a reconnaissance traffic signature?
    a. Trojan program
    b. ping sweep
    c. denial of service
    d. Ping of Death

13. Which program keeps track of services and ports made available through Remote Procedure Calls?
    a. Network Information System
    b. Network File System
    c. Network File Sharing
    d. portmapper

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

14. To avoid attacks that use advanced evasion techniques, such as path obfuscation, CGI scripts, and packet injection, you must do which of the following? (Choose all that apply.)
    a. Watch your log files closely.
    b. Install additional IDPS sensors.
    c. Keep your anti-adware software updated.
    d. Keep your IDPS signature files updated.
15. Which of the following features distinguishes IPv6 from IPv4?
    a. IPv6 fragmentation occurs on IPv6-compliant routers.
    b. IPv4 is unfragmentable.
    c. IPv6 fragmentation occurs only at the source node.
    d. The IPv4 maximum fragment size is larger than its IPv6 counterpart.

Hands-On Projects

Hands-On Project 3-1: Setting the PATH Variable to Work with TShark

Time Required: 10 minutes

Objective: Enable TShark to run from a command prompt regardless of which directory you use.

Description: TShark is a command-line tool that comes with Wireshark. You can run TShark only from the directory in which the tshark.exe program is located, unless you modify the value of the PATH environmental variable. In this project, you modify the PATH value to accommodate TShark.

Before performing any of the following projects in this chapter, disable your Windows and ZoneAlarm firewalls.

1. Log on to either the Windows 7 or Windows Server 2008 computer. You can perform this project on either system.
2. Open a command prompt and note the current directory listed in the prompt. Type **tshark** and press **Enter**. The system is unable to find the program. Type **cd ** and press **Enter** to change to the root of the C drive, as indicated by the prompt. Type **tshark** and press **Enter**. Again, the system cannot find TShark. The operating system uses an environmental variable called PATH to determine which directories it should search for executable files like TShark. The value of this variable is stored in RAM and is easy to view. Type **PATH** and press **Enter**. A list of directories appears, separated by semicolons. Because the directory where tshark.exe is located, C:\Program

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Files\Wireshark, is not in the path, you must either change to the Wireshark directory or modify the value of the PATH variable to be able to run the TShark program. You will modify the PATH variable's value.

3. Type **PATH=%PATH%;C:\Program Files\Wireshark** and press **Enter**. Type **tshark** and press **Enter**. TShark starts and begins to capture network traffic if it identifies the correct network interface that is connected to the network. Press the **Ctrl** and **C** keys simultaneously to stop TShark. (From this point, the key combination is written as Ctrl+C.) Type **PATH** and press **Enter**. Notice that the Wireshark directory is now included in the PATH value.
4. Close the command prompt and then open a new command prompt. Type **PATH** and press **Enter**. It is evident that your alteration of the PATH value was limited to the specific shell in which you made the alteration. You must make a more permanent change to the shell.
5. Click **Start**, right-click **Computer**, click **Properties**, click **Advanced system settings**, and click **Environment Variables**. In the System variables section, scroll if needed, and click the entry that starts with the word **Path**. Click **Edit**. Click in the Variable value text box without altering the current value. Press **End**, and then type **;C:\Program Files\Wireshark**. Click **OK** three times, and close the System window.
6. Return to the command prompt, and type **tshark**. Regardless of which directory you are in, the TShark program now runs. Press **Ctrl+C** to stop TShark.
7. If you see a message that prompts you to run the diskperf -y command to enable disk statistics, you should follow those directions. Type **diskperf -y** and press **Enter**.
8. Leave your system running for the next project.

Hands-On Project 3-2: Using TShark to Capture ICMP Packets

Time Required: 15 minutes

Objective: Capture ICMP packets using TShark.

Description: You now use TShark to capture ICMP packets. You also start the Wireshark help utility to explore and use some of the TShark options. Finally, you save the TShark captures to a file instead of viewing them live as they are captured, which allows you to analyze captured packets and keep a record of them.

1. From a command prompt, type **user-guide.chm** and press **Enter**. The Wireshark User's Guide appears. Expand the guide by clicking the + button to its left. Expand the **Related command line tools** node, click **tshark: Terminal-based Wireshark**, and review the information.
2. Arrange your desktop so that two command prompt windows are visible. In the first command prompt window, type **tshark -O icmp -V -x > C:\test1.txt**, but do not press Enter. In the second command prompt window, type **ping *partner's_IPv4_address***, where *partner's_IPv4_address* is the 4 octets of your partner's IP address. See Figure 3-17 for an example.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

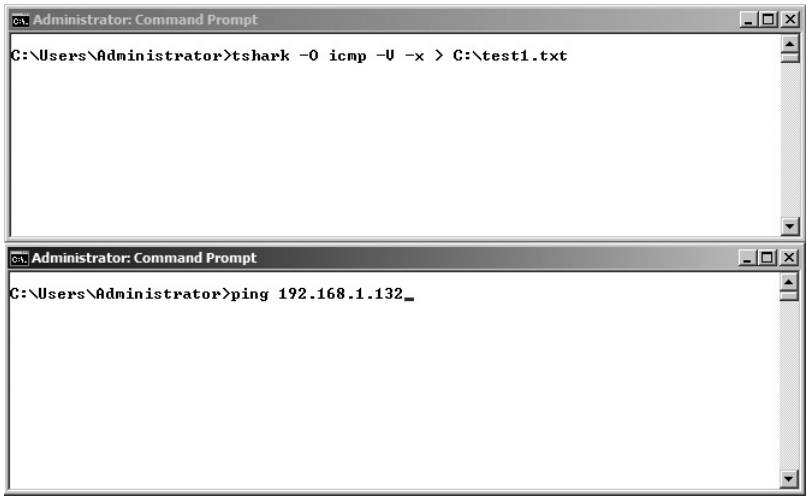

Source: Used with permission of Microsoft Corporation

Figure 3-17 Preparing a TShark ping capture

3. Review the TShark help page to examine the functions of the options you used in the TShark command in Step 2. Click the first command prompt window, and then press **Enter** to run the TShark command.

If you receive an access-denied error at the command prompt, you need to run the command prompt as an administrator. Click **Start**, click **All Programs**, click **Accessories**, right-click **Command Prompt**, and click **Run as administrator**.

Click the second window, and press **Enter** to begin the ping. When the ping is complete, click the first command prompt window, and press **Ctrl+C** to stop TShark. Notice that no capture output is displayed in the TShark command prompt window because the output was redirected from the screen by the ">" option. Instead, the output was saved as a text file to C:\test1. The "-O icmp" option indicates that only packets using the ICMP protocol are captured. The "-V" option means that a verbose capture will be run, and the "-x" option means that hexadecimal and ASCII data should also be captured.

4. From either command prompt window, type **notepad C:\test1** and press **Enter**. The captured packets appear in a Notepad document. Figure 3-18 shows one of the ICMP packets captured by TShark. You should be able to identify all the elements of this capture.

5. Using the "-w" option in TShark, you can save captures in a binary format that Wireshark's graphical user interface can open. Leave your systems running for the next project.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

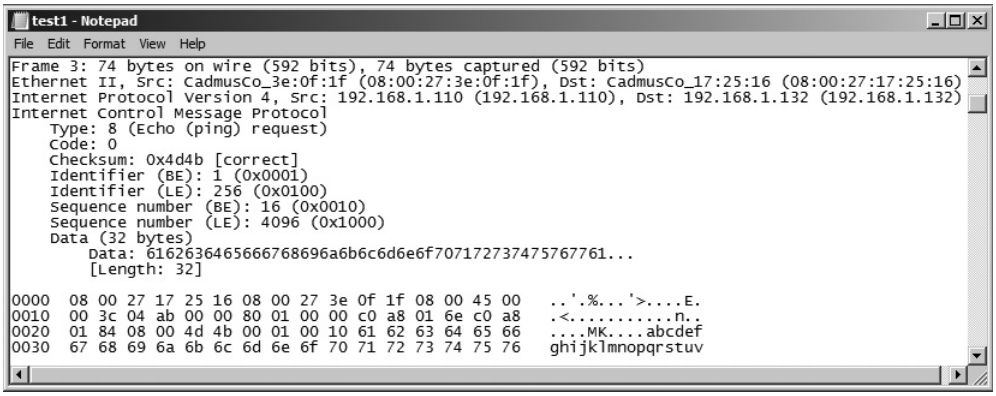

Source: Used with permission of Microsoft Corporation

Figure 3-18 A TShark ping capture

Hands-On Project 3-3: Installing an FTP Server

Time Required: 10 minutes

Objective: Install the FTP service on Windows Server 2008.

Description: To explore normal FTP traffic, you install the FTP service on your Windows Server 2008 computer.

1. Access the Windows Server 2008 computer with an administrative account.
2. Click **Start**, click **Administrative Tools**, and click **Server Manager**. In the left frame, click **Roles**. In the Roles frame, click **Add Roles**. In the Before You Begin window, read the message, click the box that tells Windows to skip this page by default, and click **Next**. In the Select Server Roles window, click the **Web Server (IIS)** check box, and then click **Next**. In the Web Server (IIS) window, read the message and click **Next**. In the Select Role Services window, scroll down, click the **FTP Server** check box, and click **Next**. In the Confirm Installation Selections window, click **Install**. When the Installation succeeded message appears, click **Close**.
3. In Server Manager, expand the **Roles** node in the left frame. Expand the **Web Server (IIS)** node, and click **Internet Information Services (IIS) Manager**.
4. In the Connections frame, expand your server node. Your screen should look similar to Figure 3-19.

   In the TEAM*x*SRV Home frame (where *x* is the team number assigned by your instructor), double-click **FTP Authentication**. Notice that Anonymous Authentication is disabled. Many FTP servers let you log on without an account. When prompted for a username, you can type **anonymous**; for your password, you can type anything, but you should type your e-mail address. In this project, you require a user account on your domain for a user to access the FTP server.
5. In the Connections frame, click **Sites**. In the Sites frame, notice that a Web site is created by default. In the Actions frame, click **Add FTP Site**. In the Site Information window, click the **…** button in the Content Directory section. Browse to **C:\inetpub\ftproot**, select it, and

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

click **OK**. In the FTP site name text box, type **Team*x*FTP**, where *x* is the team number assigned by your instructor. Your window should look like Figure 3-20. Click **Next**.

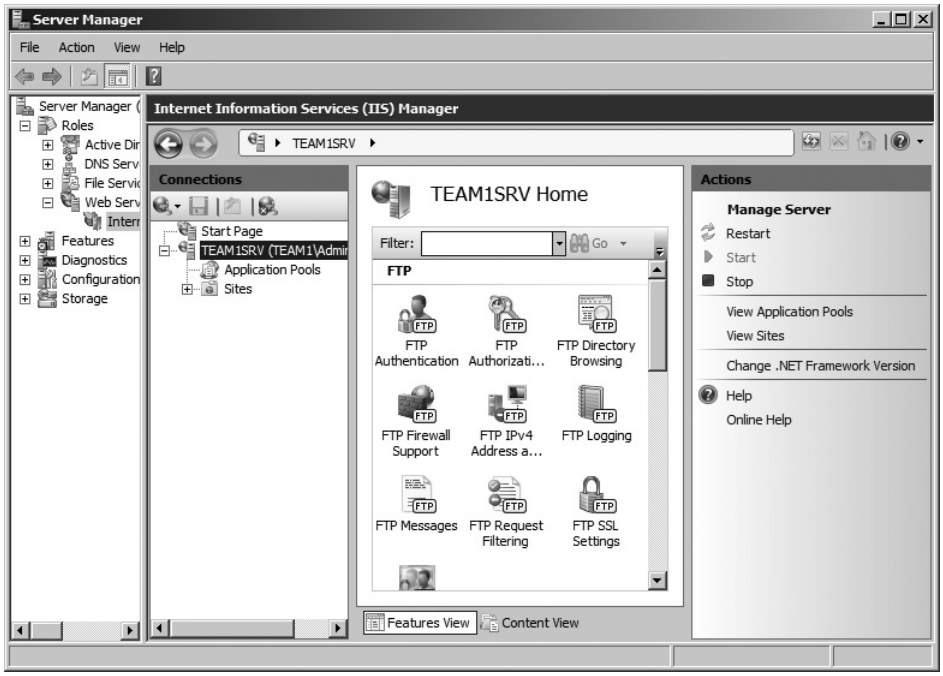

Source: Used with permission of Microsoft Corporation

Figure 3-19 Internet Information Services (IIS) Manager

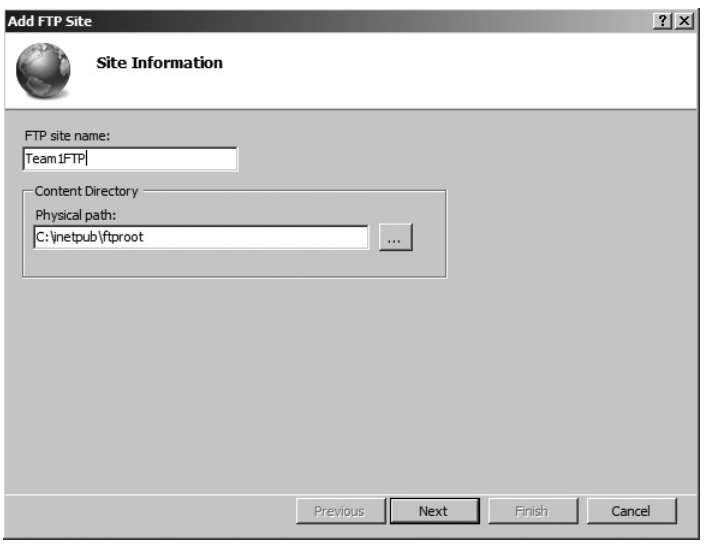

Source: Used with permission of Microsoft Corporation

Figure 3-20 FTP Site Information window

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

In the Binding and SSL Settings window, notice that your FTP server will listen at the usual FTP control port 21. The FTP server will use the usual FTP data port 20 for file transmission. In the SSL section, click the **No SSL** button, and then click **Next**. In the Authentication and Authorization Information window, click the **Basic** check box in the Authentication section, click the **Allow access to** drop-down arrow, and select **Specified users**. In the empty text box under Allow access to, type your first name. Click the **Permissions Read** and **Write** check boxes, and then click **Finish**. Leave Server Manager open.

6. You specified that a user can gain Read and Write access to the FTP server by entering your first name as a logon name, so you need to create that user. Click **Start**, click **Administrative Tools**, click **Active Directory Users and Computers**, and expand your Team*x* domain, if necessary. Right-click the **Users** folder, click **New**, and click **User**. In the window that appears, type **User** as the first name, type **One** as the last name, and type **UserOne** as the logon name. Click **Next**. In the password text box, type **Pa$$word** as the password, uncheck the **User must change password at next logon** box, and check the **Password never expires** box.

Remember that you created a local user account in Windows 7 when you first installed it, but that the accounts made in Active Directory are domain accounts. The accounts in Windows 7 are similar to TEAM1CLIENT\Dean, and the domain account you just made is similar to TEAM1\Dean. Remember, *TeamxSRV* is the hostname of the Windows Server 2008 computer, *TeamxClient* is the hostname of the Windows 7 system, and Team*x* is the name of the domain to which the computers both belong.

Click **Next** and then click **Finish**. Close the Active Directory Users and Computers window.

7. Return to Server Manager. In the Connections frame, select **Team*x*FTP**. In the Actions frame, click **Edit Permissions**. In the ftproot Properties window, click the **Security** tab, and then click **Edit**. In the Permissions for ftproot window, click **Add**. In the next window, type your first name, and click **Check Names**. Your user account should appear in the box. Click **OK**. In the Permissions for *Your_first_name* section, click the **Modify** check box in the Allow column. Click **OK** two times.
8. Double-click **FTP Authorization Rules**. In the Actions frame, click **Add Allow Rule**. In the next window, click the **Permissions Read** and **Write** check boxes. Click **OK**.
9. Close Server Manager. Click **Start**, click **Computer**, and navigate to C:\inetpub\ftproot. Right-click the white space in the right frame, click **New**, and click **Text Document**. Name the document **FTPTest**, and type **This is a test** in the body of the document. Save and close the document. Close Windows Explorer.
10. Turn off any firewall that is running on the server. Leave your server running for the next project.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Project 3-4: Capturing FTP Traffic with TShark

Time Required: 10 minutes

Objective: Connect from the Windows 7 computer to the Windows Server 2008 FTP service, and then capture and analyze the traffic produced.

Description: To become familiar with different types of traffic and learn to distinguish normal traffic from suspicious traffic, you should examine as many packet captures as you can. In this activity, you transfer a file from the FTP server to the Windows 7 client and analyze the traffic.

1. Log on to your Team*x* domain as the user you created in Step 6 of Hands-On Project 3-3.
2. Click **Start**, click **All Programs**, click **Accessories**, right-click **Command Prompt**, and click **Run as administrator**. In the User Account Control window, you are asked to provide administrative credentials for the Team*x* domain. Type **Administrator** in the User name text box, type **Pa$$word** in the Password text box, and click **Yes**. At the command prompt, type **tshark -x -n > C:\FTPcapture.txt**, but do not press Enter.
3. Open another command prompt window without running as administrator. In the second command prompt window, type **ftp Team*x*SRV.Team*x*.net**, but do not press Enter. Arrange the command prompt windows so you can see both. Record the prompt in the second command prompt window using the following space:

   _____

   You will go to this directory later to see the file you download from the FTP server.
4. Click the first command prompt window that will run TShark, and press **Enter**. Immediately click the second command prompt window that will run FTP, and press **Enter**.
5. In the FTP command prompt window, type your first name at the user prompt, and press **Enter**. At the password prompt, type **Pa$$word** and press **Enter**. When you see the User logged in message, type **ls** and press **Enter** to see a list of files in the FTP server directory. You should see the FTPTest.txt file that you created earlier.
6. Type **get FTPTest.txt** and press **Enter**. When you see the Transfer complete message, type **bye** and press **Enter** to exit FTP. In the other command prompt window, press **Ctrl+C** to stop TShark.
7. At either command prompt, type **notepad C:\FTPcapture.txt** and examine the file contents. The TShark command specified that you would capture only FTP traffic, and that hexadecimal and ASCII data would be included in the capture.
8. You should be able to identify the TCP three-way handshake, authentication, and file transfer. This type of FTP traffic is not encrypted, so you should be able to find details, including the password in the ASCII capture sections. Explain your findings to your instructor.
9. Leave your systems running for the next project.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Project 3-5: Exploring Nmap Scans

Time Required: 30 minutes

Objective: Explore signatures of Nmap scans.

Description: A network scanner is commonly used to attempt digital attacks. In this activity, you capture traffic from SYN and Xmas scans generated by Nmap to become familiar with the appearance of scanning traffic.

1. Log on to the Windows 7 and Windows Server 2008 computers with administrative accounts.
2. Start Nmap.
3. In the Command box, enter **nmap -sS *networkAddress***, where *networkAddress* is the network address of your classroom network, as provided by your instructor. For example, type the following to produce an Nmap SYN scan of your network:

   nmap -sS 192.168.1.0/24
4. Open a command prompt, type **tshark -x -n > C:\SYNscan1.txt** and press **Enter**.
5. Return to the Nmap window and click **Scan**. (If any messages from ZoneAlarm appear, click **Allow**.) When the output reports that Nmap is finished, return to the command prompt and type **Ctrl+C**. Type **notepad C:\SYNscan1.txt** and press **Enter**.

If you receive an access-denied error at the command prompt, you need to run the command prompt as an administrator. Click **Start**, click **All Programs**, click **Accessories**, right-click **Command Prompt**, and click **Run as administrator**.

6. Examine the TShark log of the scan. Nmap performs ARP scans to determine which systems are running and then performs a SYN scan on the "live" systems. Your network might have so much traffic that you could have trouble isolating the packets associated with your own computer's scan activity. To address this problem, you can create a capture filter.
7. From a command prompt, type **getmac** and press **Enter**. Note the 6-byte MAC address of your computer here:

   _____

   Type **ipconfig** and press **Enter**. Note your IPv4 address here:

   _____

   Prepare your Nmap window to run the same scan you ran in Step 3. At the command prompt, type the following command:

   TShark -f "ether host *yourMAC* || host *yourIPaddress*" -x -n > C:\SYNscan2.txt

   where *yourMAC* is the MAC address of your computer (for example, 08-00-27-17-25-16) and *yourIPaddress* is your IPv4 address. The command specifies that the capture filter (-f) only records packets that contain either your MAC address as the source or destination MAC address or your IPv4 address as the source or destination IP address.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Press **Enter**, switch to the Nmap window, and click **Scan**. When the scan is complete, return to the command prompt, press **Ctrl+C**, and examine the log file by typing **notepad C:\SYNscan2.txt**. You should now see only packets that contain your MAC address; you do not capture other classmates' traffic unless they sent a packet to you. Be sure that you become familiar with the signature of an Nmap SYN scan.

8. Repeat the same procedure to examine an Xmas scan. The Nmap command is **nmap -sX *networkAddress***. Examine the capture and see if you can find any obvious patterns in the result.
9. Log off your systems.

Case Projects

Case Project 3-1: Evading Firewalls and the IDPS

You have been assigned to use Nmap to develop testing procedures for how attackers might try to evade detection at your organization's firewalls and IDPS. Prepare a report that explains which Nmap options can be used to evade firewalls and IDPSs, and how these options function.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 4

Routing Fundamentals

After reading this chapter and completing the exercises, you will be able to:

- Describe the routing process
- Configure a Cisco router
- Describe router security controls
- Create access control lists

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Routers are among the most common network infrastructure devices found on local area networks (LANs) and wide area networks (WANs). Because they handle all types of network traffic, they are important targets for hackers, who want to learn about your network as they build an attack on your servers and workstations. For example, DNS zone transfers, SNMP (Simple Network Management Protocol) network status information, and information about the network's routes and subnets all travel through routers and switches.

Not only are routers and switches exposed to malicious attacks, they can significantly limit data availability if they are misconfigured or malfunction. While firewalls, intrusion detection and prevention systems (IDPSs), and proxy servers may seem more glamorous, you need to pay close attention to the security and functional ability of routers to optimize network security.

This chapter begins with an examination of the routing process and follows the delivery of packets from the perspective of the Network and Data Link layers. You learn about the router function, including router processes such as routing tables and protocols, and learn how to connect to and configure a router. Next, you examine router security controls, including access control lists, authentication, authorization, logging, and the use of encrypted protocols to secure router configuration connections.

Examining the Routing Process

Routing is the process of transporting packets of information across a network from the source node to the destination node. Routers determine the best path for the packet to take and then send the packet on its way. To determine the best path, routers use metrics such as hop count, bandwidth, or link state. However, the administrator can also configure predetermined paths for packets based on the protocols being transported and other variables. Routing takes place at the Network layer of the OSI (Open Systems Interconnection model).

While the routing process can be complicated, fundamentally the following steps take place:

1. An application, utility, or service on a source computer generates a packet to send to a specific destination. The OSI Network layer function of the source computer determines whether the destination is on the same network segment as the source computer or on a different one. If the destination is on the same network as the source, the packet is sent directly. If the destination is on a different network from the source, the packet is sent to the interface configured on the source as the default gateway—the interface on a router that gives a computer access outside its own network.
2. The router receives the packet and strips off the Data Link layer header, which includes source and destination MAC (media access control) addresses. Next, the router examines the destination Network layer address.
3. Based on the destination address and the information the router maintains in its routing table, the router determines which of its multiple interfaces to use to move the packet closer to its destination.
4. The router builds a header that is appropriate for the selected outbound interface. For example, if the outbound interface uses a token ring network, the router builds a token ring header so that the interface can understand the information and process the packet.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The header contains Data Link layer source and destination addresses and protocol information.

5. The packet is sent through the selected interface to the next hop.

This explanation of routing has been simplified. Typically, each packet requires the router to perform additional processing tasks, such as checking access control lists to filter traffic or tunneling different protocols. For example, an AppleTalk or IPX/SPX packet might be sent through an IP network. Generally, protocols do not "talk" to each other without an intermediary device to provide translation.

The Address Resolution Protocol Processes

Regardless of whether you use IP, IPX, or another Network layer protocol, a packet cannot reach its destination until the Data Link layer MAC address is determined for the destination and the Network layer address is resolved to the MAC address. **Address Resolution Protocol (ARP)** resolves IP addresses to MAC addresses, and is one of the subprotocols of the TCP/IP networking protocol stack.

Many network devices, including workstations, maintain their own **ARP tables** that list the MAC and IP address resolutions of other devices on the network. These ARP tables may include static entries that remain in the table and dynamic entries that are populated during the ARP resolution process explained below. These dynamic entries have a limited time to live. For example, in Windows workstations, a dynamic ARP table entry is automatically removed after two minutes.

When a computer prepares to transmit a packet to a destination on the local network, it checks its ARP table for an IP-to-MAC address resolution for the destination node. If the computer finds the address resolution, the source computer uses the information to create the Data Link header with the source and destination MAC addresses. The packet is then sent directly to the destination node on the local network.

If the source computer does not find an entry for the destination computer's IP address in its ARP table, it sends an ARP broadcast to the local subnet in an attempt to discover the destination computer's MAC address. (See Figure 2-15 for a packet capture of this process.) Because the destination address in this discovery packet is a broadcast (FF-FF-FF-FF-FF-FF), every host on the local subnet must process the packet to determine if it is "of interest." When the packet reaches the Network layer, where ARP is processed, the host can determine if its own IP address is being specified in the ARP broadcast packet. Only the correct host responds; the others discard the packet. The host that discovers its IP address in the broadcast packet responds to the source computer's ARP request by providing its MAC address.

The source computer updates its ARP table with the destination computer's MAC address, creates and addresses the data packet, and sends it.

Remember from Chapter 2 that ARP is not needed in IPv6 because Neighbor Discovery provides IP-to-MAC resolution.

Accessing a Router

To understand router commands and configurations, you need to know how to work at the **Cisco Internetwork Operating System (IOS)** command-line interface. Before you can do that, however, you must gain access to the router. Several methods are commonly used.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The back of a Cisco router contains several interfaces—network modules such as Ethernet, Frame Relay, and Fiber Channel, as well as the power switch and other devices specific to the router model. The **auxiliary (AUX) port** and **console (CON) port** are particularly important when you access the command line for configuration, troubleshooting, and maintenance. When you have physical access to the router, you can use the CON port with a rollover cable attached to a laptop or other workstation. A **rollover cable** is similar to unshielded twisted pair (UTP) cable, which is commonly used to connect Ethernet devices. However, the arrangement of the connectors is different: Pins 1 through 8 on one end of the cable connect to pins 8 through 1 on the other end. The Ethernet jack on one end of the rollover cable is connected to the router's CON port, and the other end is connected to the management computer's COM port, or to an RJ-45 to DB-9 adapter that connects to the management computer's COM port.

Again, remember that ARP is not needed in IPv6 because Neighbor Discovery provides IP-to-MAC resolution.

Routing Tables

Routing tables are lists of networks that contain information for reaching the networks, and indicators such as hop count and link-state that help the router determine the most efficient route. Routing tables have three types of entries. **Static routes** are entered in a routing table by an administrator. **Dynamic routes** are populated automatically by routing protocols and routing algorithms that the router uses to calculate the best path. The information used to create dynamic routes is obtained from other routers. **Default routes** are manually configured routes that direct all packets not specifically configured in the routing table. A packet that cannot be routed to its destination via learned or static routes is sent to the interface specified in the default route automatically.

Cisco routers use three main processes when building and maintaining routing tables:

- Routing processes that actually run a routing protocol such as Interior Gateway Routing Protocol or Open Shortest Path First
- The forwarding process, which requests information from the routing table for making forwarding decisions
- Routing tables from other routers that are sent in response to requests for information about routing and forwarding processes, or that are sent automatically as default updates depending on the routing protocol used

Routing protocols are discussed in more detail later in this chapter.

Static Routing

Routing protocols are used for communication between routers and to support dynamic entries in the routers' routing tables. In some circumstances, dynamic routes are not necessary. Routing protocols use network bandwidth and consume resources; because they contain information about the network, they are also a security concern. If the network can run efficiently using only static routes, dynamic routes can be eliminated. An example of a situation in which only static routing is needed is a **stub network** (see Figure 4-1). A router with only one route is called a **stub router**. A stub router is usually found at the end of the

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

network line and is connected to only one other router. Stub networks are generally found at the network's edge and are considered dead-end segments.

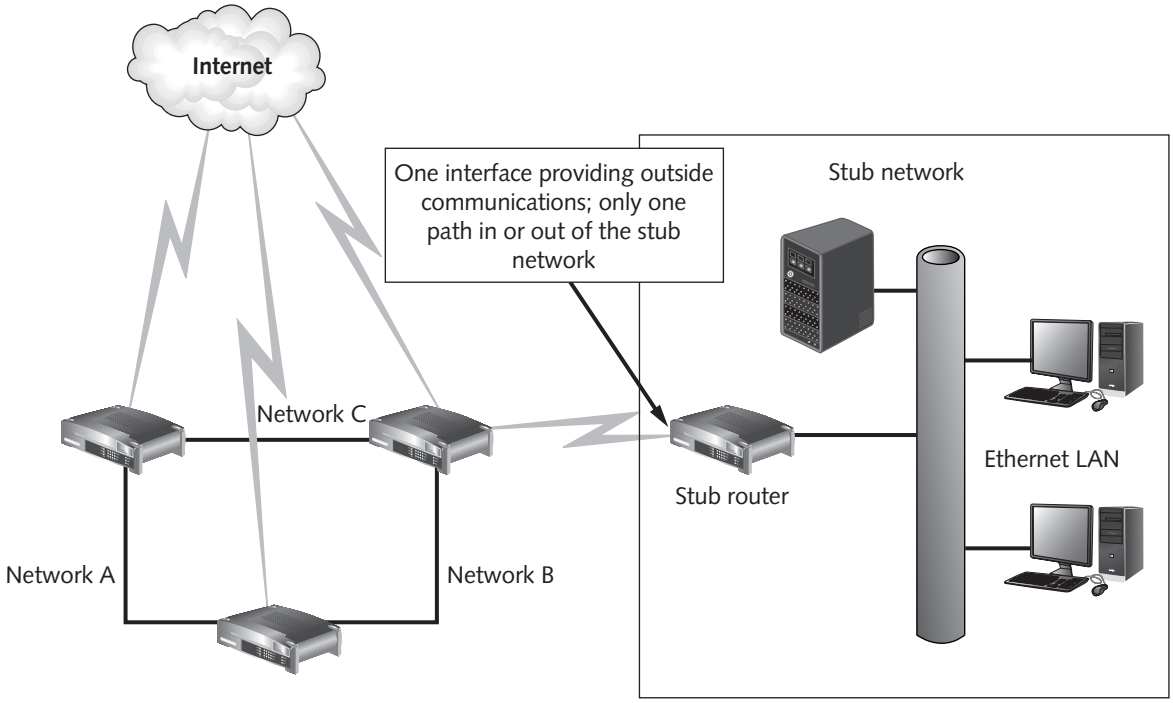

© Cengage Learning 2014

Figure 4-1 Stub network

At times, an administrator might need to specify certain routes or adjust traffic flow to maximize efficiency, improve security or performance, and conserve bandwidth. In these cases, the administrator can enter the route manually in the routing table. A static route, by default, has a lower metric than one determined by a dynamic routing protocol. Routers forward traffic based on metrics; when more than one route can send a packet to its destination, the router chooses the route with the lowest metric. Static routes are configured on Cisco routers using the ip route command, which has the following syntax:

ip route [*destination network*] [*destination network subnet mask*] [*IP address of the next hop interface*] [*administrative distance*]

The main disadvantages of static routing are the time required to configure routes and the effort needed to maintain the routing table when the network's topology changes. The advantages of static routing are the fine-tuned control that administrators have over traffic flow as well as a significant decrease in resource use. When using static routing, the router does not need to expend resources deciding which route to use. In addition, routers do not generate network traffic for route discovery because the static route is configured manually, thus reducing the overall traffic load on the network.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Dynamic Routing

Routing protocols use algorithms that enable routers to communicate with each other and map the network in the form of routing tables. Routing tables are updated at regular intervals or when a route changes, depending on the protocol. While static routing can be appropriate in small networks, as networks grow they become more complicated and any changes in routes or route accessibility require administrative effort to maintain **convergence**—the state in which all network routers have up-to-date information about the network topology. Static routing is not appropriate in larger networks. When configured to use dynamic routing protocols, the routers communicate with each other, providing information on new routes and routes that have gone down or that have returned to service, as well as providing dynamically updated metrics to determine the most efficient route.

Routing protocols, whether static or dynamic, are different from routed protocols such as TCP/IP. A **routed protocol** provides addressing for a packet being transported; a **routing protocol** allows routers to maintain information about network links in their routing tables. The two main types of routing protocols are discussed in the following sections.

Distance-Vector Routing Protocols

A **distance-vector routing protocol** uses mathematical calculations to compare routes based on some measurement of distance, such as hops. This protocol requires routers to send full or partial routing table updates periodically to neighboring routers.

Link-State Routing Protocols

A **link-state routing protocol** requires each router to maintain at least a partial network map. Routers monitor link status, and when the topology or link state changes, updates are sent to neighboring routers to inform them of the change. Link-state routing protocols use a notification called a link-state advertisement to broadcast changes. Because link-state routing protocols do not use the default periodic updates every 30 or 90 seconds, less network traffic is generated. However, these protocols do require slightly more router memory.

Routing Metrics

As mentioned earlier, routing protocols use metrics to determine the best path for packet routing. **Metrics** are cost values that help routers assess the desirability of a link. Common metrics include hop count, load, bandwidth, delay, and reliability of links.

The idea of a link "cost" does not involve money; it is a method of assigning preference ratings to a route. For example, if you have a cable link with a 10-Mbps bandwidth and a redundant DSL link with a 3-Mbps bandwidth, a router running a distance-vector routing protocol could assign an equal value to each link. The router views both links as equally desirable, even though the cable link is clearly preferable to the DSL link.

Because distance-vector routing protocols such as RIP (Routing Information Protocol) use only hop count to calculate the best path, the assessment process is prone to errors. Link-state routing protocols use multiple metrics, such as reliability and bandwidth, for path determination.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Choosing a Routing Protocol

The most common LAN routing protocols are RIP, EIGRP (Enhanced Interior Gateway Routing Protocol), OSPF (Open Shortest Path First), and IS-IS (Intermediate System-to-Intermediate System). Determining which protocol is best for your LAN depends on a number of factors:

- Administrative cost of management
- Administrative cost of configuration
- Bandwidth usage, both in the baseline state and when responding to a network event
- Frequency of network failures
- Network recovery time
- Convergence time
- Network topology

There is usually a trade-off between these factors; to accommodate one factor, you must accept deficiencies in another. The topology of the network is an important consideration because it affects the convergence time of different protocols. Network engineers must consider routing protocol selection when designing network function and topology.

Route Summarization

Route summarization (also called **supernetting**) allows service providers to assign addresses in a classless fashion and make more efficient use of available Internet addresses. IPv6 will eventually provide many more IP addresses, but IPv4 will not be phased out overnight. Until then, better ways to use IPv4 addresses are needed.

For example, to provide an organization with 2000 IP addresses, eight traditional Class C networks must be used (see Table 4-1).

Traditional Class C address range	Subnet mask	CIDR prefix
194.28.0.1–194.28.0.254	255.255.255.0	/24
194.28.1.1–194.28.1.254	255.255.255.0	/24
194.28.2.1–194.28.2.254	255.255.255.0	/24
194.28.3.1–194.28.3.254	255.255.255.0	/24
194.28.4.1–194.28.4.254	255.255.255.0	/24
194.28.5.1–194.28.5.254	255.255.255.0	/24
194.28.6.1–194.28.6.254	255.255.255.0	/24
194.28.7.1–194.28.7.254	255.255.255.0	/24

© Cengage Learning 2014

Table 4-1 Eight traditional Class C networks

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Table 4-1 lists both the subnet mask and Classless Interdomain Routing (CIDR) notation. Both show that the first three octets (24 bits) of the IP addresses represent the network identifier and that the last octet represents the host identifier. To provide packet-forwarding services to all these networks, the routers would need to have each route represented individually in their routing tables. Currently, Internet routers contain about 400,000 routes.

You can follow the daily changes of Internet routes at *www.cidr-report.org/as2.0/#General_Status*.

Route summarization allows you to create a single routing table entry that would represent all these routes. To summarize routes, convert the IP addresses to binary and then count the number of bits that are common to all networks. Table 4-2 demonstrates this process.

Class C network	Binary representation (common network bits in bold)
194.28.0.*x*	**11000010.00011100.00000**000.*x*
194.28.1.*x*	**11000010.00011100.00000**001.*x*
194.28.2.*x*	**11000010.00011100.00000**010.*x*
194.28.3.*x*	**11000010.00011100.00000**011.*x*
194.28.4.*x*	**11000010.00011100.00000**100.*x*
194.28.5.*x*	**11000010.00011100.00000**101.*x*
194.28.6.*x*	**11000010.00011100.00000**110.*x*
194.28.7.*x*	**11000010.00011100.00000**111.*x*

© Cengage Learning 2014

Table 4-2 Determination of matching network bits in each Class C network

Because all eight networks have the same first 21 bits, a single entry in a routing table for 194.28.0.0/21 would summarize all eight networks. If the first 21 bits of a packet's destination address matched the first 21 bits of a network address listed in Table 4-2, the packet would be sent to a predetermined router where the Internet meets your network: your border router. Thus, while your network routers could easily handle these eight networks, the Internet routers would only need to know one route to reach them all. With route summarization, Internet routers need about half the routing table entries (currently about 220,000) compared to all the networks to which packets can be delivered.

A concept related to route summarization is **variable length subnet masking** (**VLSM**). VLSM uses subnet masks of different lengths on the same network to assign network addresses based on need instead of using a generic masking scheme. For example, you might have a Class C network divided into subnets, with each one supporting 62 hosts. In reality, only 15 to 20 hosts are attached to each subnet, so the additional addresses are wasted. With VLSM, you can divide the network into subnets of varying sizes to support your users but make better use of your available addresses, instead of using the network-wide classful routing.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

This can be especially useful when setting the endpoint addresses for links between branch offices. Instead of using 62 host addresses to provide the subnet between the branch routers—a subnet in which only two addresses are needed—you can use VLSM to allocate only two addresses to the branch-to-branch link and use the other 60 addresses to divide among other subnets.

IPv6 Routing As you learned in Chapter 2, IPv6 is gradually replacing IPv4. Microsoft operating systems now ship with IPv6 enabled by default. The routing protocols have been updated: RIP has upgraded to the IPv6-compliant RIPng. Similarly, OSPFv3, EIGRP for IPv6, and IS-IS for IPv6 are all IPv6 compliant. The U.S. government has mandated that all of its agencies must deploy IPv6 on their public Web sites before September 30, 2012, and that they must upgrade their entire internal infrastructure to IPv6 before September 30, 2014.

Cisco routers and others are now capable of supporting IPv6. The following example examines the configuration of the Branch06 router in Figure 4-2.

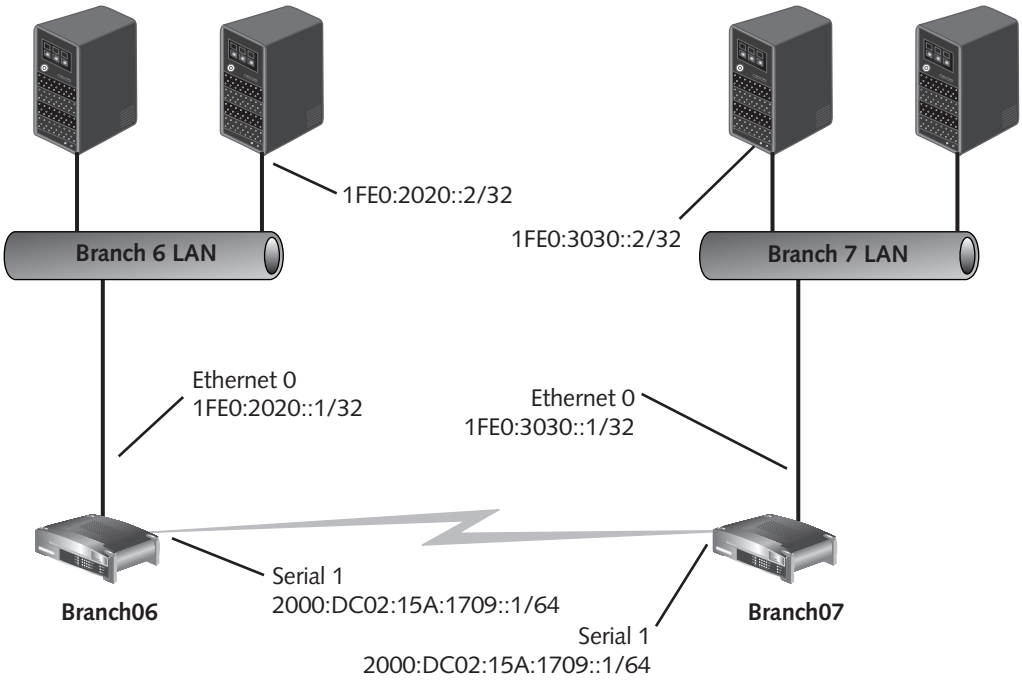

© Cengage Learning 2014

Figure 4-2 IPv6 addressing in branch networks

The Branch06 router connects the Branch 6 LAN to the Branch 7 LAN. The hosts on the Branch 6 LAN are in the IPv6 network with the prefix of 1FE0:2020, as indicated by the 32-bit mask. The addressed host has a link-local (private) IP address of 1FE0:2020::2/32, and the Branch06 router has an Ethernet 0 address on the same network, 1FE0:2020::1/32. The Branch06 router also has a connection through Serial 1 to the router at the Branch 7 LAN. This WAN network uses global unicast addresses (public) with the 64-bit prefix of 2000:DC02:15A:1709. The Branch06 router's WAN address is 2000:DC02:15A:1709::1/64.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

You would take the following steps to address the Branch06 router. First, you would enable IPv6 on the router using the following command:

```
Branch06(config)#ipv6 unicast-routing
```

The following commands configure Ethernet 0:

```
Branch06(config)#interface FastEthernet 0/0
Branch06(config-if)#ipv6 address 1FE0:2020::1/32
Branch06(config-if)#no shutdown
Branch06(config-if)#exit
```

The following commands configure Serial 1:

```
Branch06(config)#interface Serial 0/1
Branch06(config-if)#ipv6 address 2000:DC02:15A:1709::1/64
Branch06(config-if)#no shutdown
Branch06(config-if)#exit
```

After this, you could easily install the dynamic routing protocol RIPng because no more "network" statements would be required, as in RIPv1 and v2.

```
Branch06(config)#ipv6 router rip RIPng
```

The tag "RIPng" can be any combination of letters and numbers; the tag is used to identify the RIP process. Then RIP is applied to each interface that will participate in RIP routing.

```
Branch06(config)#interface FastEthernet 0/0
Branch06(config-if)#ipv6 rip RIPng enable
Branch06(config-if)#exit
Branch06(config)#interface Serial 0/1
Branch06(config-if)#ipv6 rip RIPng enable
Branch06(config-if)#exit
```

Router Security Fundamentals

Because routers contain detailed information about network topology and addressing, they are a target for malicious attacks. Furthermore, one of the three main goals of information security—data availability—can be seriously compromised when routers do not perform their duties properly. Router security is an important component of network defense. Routers also play a role in defensive countermeasures when they work in conjunction with an IDPS to block packets from a source that the IDPS has determined to be a threat.

In this section, you learn about the function and configuration of access control lists used to perform packet filtering. You examine router logging and its importance in network security. Later in this chapter, you learn about router authentication, banners, and the use of Secure Shell to encrypt configuration connections.

Creating and Using Access Control Lists

Router **access control lists** (**ACLs**) are permit or deny statements that filter traffic based on the source and destination address, source or destination port number, and protocol in the packet

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

header. ACLs provide traffic-flow control and enhance network security. They can also be used to fine-tune performance and control client access to sensitive network segments.

Use and Rules

You can approach ACLs in two ways: specify what traffic to deny and permit all other traffic, or specify what to permit and deny all other traffic. The latter approach is more secure, but it requires more planning and a more complex list. You must consider two factors when configuring ACLs:

- ACLs end with an implicit "deny any" statement, which means any packet that does not match the requirements for passage is blocked. Although this approach is more secure—and is the default behavior on Cisco ACLs—it might not serve the network's needs and could block desirable traffic. To remedy this problem in networks when you do not need to block all access, include a "permit any" statement at the end of an ACL.
- ACLs are processed in sequential order. To avoid unnecessary use of router processing resources, rules that match common network traffic should be placed higher on the list so that they are processed first and the packet does not need to be compared to a large number of rules before the appropriate rule is identified for the packet. Careful planning is necessary to ensure that allowed packets are not blocked inadvertently, and packets that should be blocked do not slip through.

Table 4-3 summarizes some of the problems you should be aware of when creating ACLs and their solutions.

Problem	Solution
Lack of planning results in simple logic mistakes.	Plan carefully what needs to be filtered and what needs access.
Sequential processing results in filtering errors.	Use the *IP Access List Entry Sequence Numbering* feature in Cisco IOS versions 12.2 and above, which allows you to move and insert rules in an existing ACL.
Applying ACLs via Telnet can result in lost connectivity for the administrator applying the list.	Use the reload command to restore access as long as the running configuration was not copied to the startup configuration.

© Cengage Learning 2014

Table 4-3 ACLs: Common problems and solutions

Remember the following general rules for ACLs:

- Routers apply lists sequentially.
- Packets are processed only until a match is made, and then they are allowed or denied.
- Lists always end with an implicit "deny any" statement.
- ACLs must be applied to an interface as inbound or outbound filters.
- The terms *inbound* and *outbound* refer to the perspective of the router; a packet entering the router is considered inbound, and a packet exiting the router is considered outbound.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- ACLs are not active until they are applied to an interface.
- Only one ACL per protocol and per direction can be applied to an interface.
- ACLs take effect immediately, but if you want the list to be permanent, you must copy the running configuration to the startup configuration using the copy running-config startup-config command.

Test ACLs thoroughly before applying them to a production router in your network; make sure they work as intended and correct any errors before you apply them. You should have a baseline for your network so that you know what "normal" traffic looks like. If you have a baseline before and after applying an ACL, you can determine where problems might occur. Ideally, you should have a baseline in your test network first.

Standard ACLs

Standard ACLs have minimal configuration options. They can filter only on source IP address information, such as a host, subnet, or network address. Of course, there is no need to create a complex ACL if a simple one will do the job. Increasing complexity leads to increased chances for errors, so you should use extended ACLs only if you have a good reason for that level of control. Standard ACLs can be configured for IP, Internetwork Packet Exchange (IPX), AppleTalk, and other protocols. The most commonly used ACLs are for IP, so the next sections focus on it.

After configuring a standard ACL, you must apply it to the interface and specify the direction of the filter. ACLs are applied to inbound or outbound packets, and only one ACL per direction can be applied to an interface at a time.

Standard IP ACLs

With standard IP ACLs, you can permit or deny traffic from a source host, a subnet, or an IP address. Destination addressing does not affect standard ACLs, which work well for simple packet filtering. Standard ACLs, like all varieties of ACLs, use an **inverse mask** that tells the router which bits in the address to be filtered are significant. (An inverse mask is indicated by the "source wildcard mask" parameter in commands.) A 0 bit in an inverse mask means to check the corresponding bit value in the IP address, and a 1 bit means to ignore the corresponding bit value in the IP address. An inverse mask of 0.0.0.0 means that all the bits are significant, so that specific IP address is filtered. To indicate the filtering of an entire subnet, mask only the host ID portion of the address. For example, the inverse mask to filter the entire 172.16.0.0 network, which has a default subnet mask, would be 0.0.255.255. The 0s are significant; the 255s are ignored.

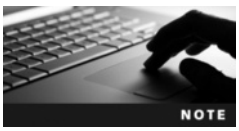

Remember that a 0 or a 1-bit value in an IP address refers to its binary format. For instance, the inverse mask 0.0.0.255 (decimal format) is the same as 00000000.00000000.00000000.11111111 (binary format).

Standard ACLs have the following characteristics:

- They can filter based on source address.
- They can filter by host, subnet, or network address using an inverse mask.
- They should be placed on the router interface as close to the destination as possible.
- They have a default inverse mask of 0.0.0.0.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Standard ACLs use the following syntax:

access-list [*list#*] [permit|deny] [*source IP address*] [*source wildcard mask*]

The following list explains the parameters in this syntax:

- *list#*—Standard IP ACLs are represented by a number from 1 to 99.
- permit|deny—Specifies a permit or deny action to be taken when a packet is identified that meets the other filtering parameters.
- *source IP address*—Indicates the IP address of the source in the packet header to be identified for filtering.
- *source wildcard mask*—Determines which bits of the source address must match for the packet to be identified for filtering.

Extended ACLs

Extended ACLs offer many more filtering options than a standard list. A standard list can filter based on source addressing, but an extended list can provide fine-tuned control over source and destination addresses, ports, and protocols that you want to filter. Of course, increased complexity means more chances of making a mistake, so be careful when creating and using extended ACLs. As with standard ACLs, extended ACLs are typically used in TCP/IP networks. However, you may need to set ACLs for IPX or other protocols.

Extended IP ACLs An extended IP ACL can filter based on source and destination address, port number, and protocol type. It uses the following syntax:

access-list [*list#*] [permit|deny] [*protocol*] [*source IP address*] [*source wildcard mask*] [*operator*] [*port*] [*destination IP address*] [*destination wildcard mask*] [*operator*] [*port*] [log]

The following list explains the parameters in this syntax:

- *list#*—Extended IP ACLs are represented by a number from 100 to 199.
- permit|deny—Specifies a permit or deny action to be taken when a packet is identified that meets the other filtering parameters; ACLs can have many lines specifying that a certain type of traffic is permitted or denied access.
- *protocol*—The IP protocol to be filtered; IP includes all protocols in the TCP/IP stack or a specific protocol, such as TCP (Transmission Control Protocol), UDP (User Datagram Protocol), IGMP (Internet Group Message Protocol), ESP (Encapsulating Security Payload), AHP (Authentication Header Protocol), GRE (Generic Routing Encapsulation Protocol), or ICMP (Internet Control Message Protocol).
- *source IP address*—Indicates the IP address of the source in the packet header to be identified for filtering.
- *source wildcard mask*—Determines which bits of the source address must match for the packet to be identified for filtering.
- *destination IP address*—Indicates the IP address of the destination in the packet header to be identified for filtering.
- *destination wildcard mask*—Determines which bits of the destination address must match for the packet to be identified for filtering.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *operator*—Less than (lt), greater than (gt), or equal to (eq); operators are used if the ACL filters a port number or range of port numbers.
- *port*—The source or destination port number of the protocol, depending on the position of the port specification in the rule.
- *log*—Turns on logging of ACL activity.

You should remember the following points about extended IP ACLs:

- Extended IP ACLs do not have a default inverse mask of 0.0.0.0. An inverse mask must be specified for the source. (Keywords such as "host" can be used.)
- Extended IP ACLs should be applied to an interface as close to the traffic source as possible.
- The "established" parameter can be used to allow incoming traffic that responds to an internal request. For example, to allow the network 101.0.0.0 to receive DNS resolution or responses to other network service requests, enter the following line in the ACL:

  `access-list 100 permit tcp any 101.0.0.0 0.255.255.255 established`

- Extended IP ACLs, like all ACLs, must be applied to an interface to be active. Also, remember that only one ACL per interface per direction can be active.
- There must be at least one permit access control entry in every access control list.

Consider the following example:

```
Branch06(config)#access-list 101 deny tcp 172.16.80.0  0.0.0.255  172.16.48.2
0.0.0.0 eq 23
Branch06(config)#access-list 101 permit ip any any
Branch06(config)#int e0
Branch06(config-if)#ip access-group 101 in
```

The first command creates an extended ACL numbered as 101 that denies any packets from the 172.16.80.0 network that are being sent to a specific host with the IP address of 172.16.48.2 if the destination port is 23. Because 23 is the Telnet port, it appears that the rule is designed to prohibit Telnet access to a specific host.

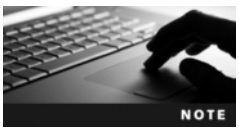

Different versions of Cisco routers may require different terms to represent interfaces. In the preceding example, the first Ethernet interface (ethernet 0) is abbreviated as "e0." Other models require the term "FastEthernet 0/0," often abbreviated as "Fa 0/0."

Named ACLs

Since IOS version 11.2, Cisco has supported **named ACLs**. Rather than referring to an ACL as a number, as in the preceding examples, you can assign a name to the ACL. Named ACLs are easier to identify, and they support more advanced features such as filtering traffic based on IP options, TCP flags, and TTL (time to live), as well as filtering noninitial fragments of packets.

Numbered and named ACLs do not use the same syntax. To name an ACL, you use the following syntax:

ip access-list [*type*] [*name*]

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The following list explains the parameters in this syntax:

- *ip access-list*—This command is used to create the list name.
- *type*—This specifies the type of list; for example, you can specify extended or standard.
- *name*—This is the name to be assigned to the ACL.

The following example shows a command to create a named ACL with the name ResearchLAN. This process begins in global configuration mode:

```
ip access-list extended ResearchLAN
```


You can see an example of using named ACLs in the "Antispoofing Logging" section later in this chapter.

Examining Cisco Router Logging

Logging is a vital component of security because it provides information for troubleshooting, monitoring traffic patterns, and discovering and tracking down possible security incidents. Cisco routers can log a variety of events, and they use the following types of logging:

- *AAA logging*—**Authentication, authorization, and accounting (AAA) logging** collects information about remote user connections, commands issued, logons, logoffs, HTTP access, and similar events. AAA logs are sent to an authentication server by using the **Terminal Access Controller Access Control System Plus (TACACS+)** protocol, the **Remote Authentication Dial-In User Service (RADIUS)** protocol, or a combination of both.
- *SNMP trap logging*—**Simple Network Management Protocol (SNMP) trap logging** sends notifications of system status changes to SNMP management stations. This logging method is normally used with an existing SNMP infrastructure.
- *System logging*—Depending on the system configuration, system logging reports system logs to different locations, including the system console port, UNIX servers via the syslog protocol, or a local logging buffer in router RAM.

Log events can also be monitored via remote sessions by using **virtual terminal (VTY)**, or TeleTypewriter or text telephone (TTY) lines. You can gain remote access to the router command line in several ways, but all inbound connections are made with TTY lines. The most important security events recorded by system logging are changes to the system configuration, ACL matches, interface status changes, and optional firewall or IDPS events.

Logging Levels

Logging events are tagged with an urgency or severity level ranging from 0 to 7, with 0 indicating the highest urgency and 7 the lowest. Table 4-4 lists these severity levels.

Logging destinations can be configured with a severity level so that events below that level are not recorded. If your router is logging all levels, your log files fill up quickly, increasing the risk of overwriting critical information. You can specify various destinations for logging, and you can buffer and view logging messages by using the show logging command at the privileged exec mode prompt.

This method has disadvantages, however. First, you must be at the terminal or access it remotely to view logs. Second, and more importantly, buffered logging is limited by the amount of memory in the router. Large log files cannot be stored in the router's memory buffer without the risk of serious performance problems and lost logs.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Level	Urgency
0	Emergency—system is unusable
1	Alert—requires immediate action
2	Critical—indicates a critical condition
3	Error—indicates an error condition
4	Warning—specifies a warning condition
5	Notification—indicates a normal but possibly significant condition
6	Informational—displays an informational message
7	Debugging—displays a debugging message

© Cengage Learning 2014

Table 4-4 Cisco router logging severity levels

Logging Options For larger log files, you can use a **syslog server**. Most Windows, UNIX, or Linux servers can be configured to host router logs. Use the logging host command to direct a router to send logs to a specific location.

Figure 4-3 shows options for the logging command. As you can see, a number of options are available for logging destinations and severity levels. You can specify the name or IP address of the logging host and set logging levels for the console, terminal lines, or syslog server.

```
Branch03(config)#logging ?
  Hostname or A.B.C.D  IP address of the logging host
  buffered             Set buffered logging parameters
  buginf               Enable buginf logging for debugging
  cns-events           Set CNS Event logging level
  console              Set console logging parameters
  count                Count every log message and timestamp last occurrence
  esm                  Set ESM filter restrictions
  exception            Limit size of exception flush output
  facility             Facility parameter for syslog messages
  filter               Specify logging filter
  history              Configure syslog history table
  host                 Set syslog server IP address and parameters
  monitor              Set terminal line (monitor) logging parameters
  on                   Enable logging to all enabled destinations
  origin-id            Add origin ID to syslog messages
  queue-limit          Set logger message queue size
  rate-limit           Set messages per second limit
  reload               Set reload logging level
  server-arp           Enable sending ARP requests for syslog servers when
                       first configured
  source-interface     Specify interface for source address in logging
                       transactions
  trap                 Set syslog server logging level
  userinfo             Enable logging of user info on privileged mode enabling
```

Source: Cisco ASA firewall

Figure 4-3 Options for the logging command

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Syslog is a simple protocol that sends a small text message via UDP or TCP to the server hosting the logs. Syslog was originally developed for UNIX and Linux systems, but variations for other systems are available. A detailed discussion is beyond the scope of this book, but you can learn more by reading RFCs 3164 and 3195.

Buffered Logging You can set parameters for **buffered logging**, which stores log output files in the router's memory (RAM). Although buffered logging has limitations, it is useful for troubleshooting purposes. Figure 4-4 shows options for the logging buffered command.


```
Branch03(config)#logging buffered ?
  <0-7>              Logging severity level
  <4096-2147483647>  Logging buffer size
  alerts             Immediate action needed           (severity=1)
  critical           Critical conditions               (severity=2)
  debugging          Debugging messages                (severity=7)
  emergencies        System is unusable                (severity=0)
  errors             Error conditions                  (severity=3)
  filtered           Enable filtered logging
  informational      Informational messages            (severity=6)
  notifications      Normal but significant conditions (severity=5)
  warnings           Warning conditions                (severity=4)
  xml                Enable logging in XML to XML logging buffer
  <cr>
```

Source: Cisco ASA firewall

Figure 4-4 Options for the logging buffered command

Antispoofing Logging **Antispoofing** is a way to prevent spoofing and ensure that no packets arrive at your security perimeter with a source address of your internal network or certain well-known or reserved addresses. Antispoofing is accomplished by using ACLs. Because ACLs are so vital to security and are the primary means of implementing basic security on a router, you learn more about them in this section.

Your ACL should instruct the router to deny any inbound packet with a source address that matches your internal network, broadcast, and loopback addresses; illegal addresses, such as all 0s or all 1s; and multicast or experimental address classes. At the end of each rule in the ACL, specify that packets matching these conditions will be logged, as shown in this example:

```
deny any 172.16.0.0 0.0.255.255 any log
```

By adding the log keyword to the end of an extended ACL, you tell the router to send information about matching packets to the router's log. When configuring the ACL, you can specify a remote logging host, as shown in the following example, in which a named ACL is created to provide antispoofing based on the parameters discussed earlier in this section:

```
ip access-list extended ResearchLAN
remark Antispoofing ACL
deny icmp any any redirect
deny ip 180.50.0.0 0.0.255.255 any log
deny igmp 224.0.0.0 31.255.255.255 any
```

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

```
permit ip any any log
exit
logging 180.50.0.12
interface FastEthernet 0/0
ip access-group ResearchLAN in
```

In this example, a named ACL, ResearchLAN, is created. A notation is made in the ACL to help identify its purpose. Attackers often use ICMP redirects to disrupt the function of a router, but the first **access control entry** (**ACE**) denies these packets. The second ACE denies and logs any packets that attempt to enter the network if they "pretend" that they come from inside the network. (This network has an address of 180.50.0.0/16.) The third ACE blocks IPv4 multicasts, and the last ACE allows all other packets; however, these packets are logged as well. The logging command specifies the IP address of the computer that hosts the log files, and the last two lines in the preceding example link the ACL to the Ethernet 0 interface, specifying that inbound packets will be filtered.

You may be wondering about the difference between the access-list and ip access-group commands. The access-list command creates the ACL. The ip access-group command assigns an ACL to an interface and specifies the direction, inbound or outbound.

Once the ACL is created and applied to an interface, you can use the show ip access-lists command from privileged exec mode to review the ACLs, as shown in Figure 4-5. Note that lines are preceded with numbers that are incremented by 10. These numbers allow you to modify or insert ACEs without having to reconstruct the entire ACL from scratch.

```
Branch03#show ip access-lists
Extended IP access list ResearchLAN
    10 deny icmp any any redirect
    20 deny ip 180.50.0.0 0.0.255.255 any log
    30 deny igmp 224.0.0.0 31.255.255.255 any
    40 permit ip any any log
```

Source: Cisco ASA firewall

Figure 4-5 Output of the show ip access-lists command

Cisco Authentication and Authorization

Authentication and authorization on a router work much like they do with a server. They identify users and allow or deny access based on the users' credentials. Authentication is the process of determining that users are who they say they are. Authorization specifies what users are allowed to do after they have accessed the system. For example, Bob is allowed to *authenticate* to the router using his username and password to check an interface's status, but he is not *authorized* to change settings on the interface.

The two types of authentication on a Cisco router are AAA and non-AAA. AAA, as mentioned previously, stands for authentication, authorization, and accounting. Cisco's AAA architecture uses one or more of three security protocols to enhance security: TACACS+, RADIUS, and Kerberos. TACACS+ is a proprietary Cisco protocol that uses TCP (port 49) for transport and encrypts all data. It also allows multiple levels of authorization and can use

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

other authentication methods. RADIUS is an open standard that uses UDP (ports 1812 and 1813) and encrypts only passwords.

Although AAA is the recommended method for access control, you can also use non-AAA methods, such as local username authentication or enable password authentication. Any method that does not use Cisco AAA Security Services is considered non-AAA. The following sections explain passwords that you can configure on Cisco routers.

Router Passwords

Passwords are the main defense against attacks on your router, but remember that if attackers have physical access to your router, they can access all its configuration settings. Cisco routers have five types of passwords you must be able to configure:

- Enable
- Enable secret
- AUX
- VTY
- Console

Because a router's main purpose is to connect networks, you might see Ethernet, token ring, Fibre Distributed Data Interface (FDDI), T1, T3, ATM (Asynchronous Transfer Mode), or other interface modules. Many Cisco routers are modular, so you can customize them by selecting the interface modules you want and specifying the encryption levels and RAM that your router must have. As mentioned, setting AAA security is the preferred method of securing access to the router; however, you still need to have local access to privileged mode in case of router or network problems.

Before setting passwords, you must know some requirements. Passwords must be 1 to 25 characters long. Leading spaces in the password are ignored, but any other spaces in it are considered part of the password. Also, the first character cannot be a number. In addition, Cisco passwords have three levels of encryption: type 0, which provides no encryption; type 7, which is encrypted but can be decrypted by router password-cracking tools that are readily available on the Internet; and the strongest level, type 5, which is a Message Digest 5 (MD5) hash. Although MD5 is a one-way hash and cannot be decrypted, it is still susceptible to brute-force attacks. Because of the relative weakness of this password system, it is important to audit the use and configuration of routers.

Enable Passwords The enable password's main purpose is to prevent casual or accidental access to privileged exec mode. Because it uses weak encryption, it provides no real security against more determined and knowledgeable attackers. Cisco recommends using an enable secret password instead of an enable password.

Enable Secret Passwords The enable secret password uses type 5 encryption and overrides an enable password. The following line shows you how to set an enable secret password:

```
Branch06(config)#enable secret Pa33m04d
```

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Remember that no password can defend against intrusion if an attacker can physically access your router; resetting a password on a router is simply too easy.

AUX Passwords Normally, the auxiliary port on a router is connected to a modem to allow remote access for router configuration. Configuring an AUX password is much like configuring a console password:

```
Branch06(config)#line aux 0
Branch06(config-line)#password M0d3m
Branch06(config-line)#login
```

For more information, see "Console Passwords" later in this section.

VTY Passwords Most Cisco routers support up to five simultaneous VTY sessions; by default, no passwords are assigned to these sessions. Cisco's built-in security on VTY lines requires configuring passwords to access the router through a VTY session, usually with Telnet. The following example shows how to set all five lines at once, although you can set each line separately if you want:

```
Branch06(config)#line vty 0 4
Branch06(config-line)#password Kodiak
Branch06(config-line)#login
```

Console Passwords You use the console port to directly connect a router to a laptop or other computer using a program such as HyperTerminal or PuTTY. The console port is normally used to set up a new router, but it can also be used to reset a router's password. For this reason, physical control of the router is critical. For security reasons, some organizations prohibit any router configuration except from the console port, eliminating the chance that transmissions can be sniffed. The following example shows how you might configure a console password:

```
Branch06(config)#line console 0
Branch06(config-line)#password J3N1ff54
Branch06(config-line)#login
Branch06(config-line)#end
```

Encrypting Passwords By default, the enable secret password is the only encrypted password type. By using the service password-encryption command in global configuration mode, as shown in the following example, you encrypt all passwords on the router using reversible encryption:

```
Branch06(config)#service password-encryption
```

You can verify the password encryption state by using the show running-configuration command from privileged exec mode. Figure 4-6 shows part of the command output on a system in which all passwords were set to "Pa$$word" but the service password-encryption command was not entered.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

It is unwise to set all passwords to be the same. Furthermore, "Pa$$word" is a weak password because it is based on a word in the dictionary, and password-cracking programs know that "$" can mean "s."

```
line con 0
 password Pa$$word
 login
line aux 0
 password Pa$$word
line vty 0 4
 password Pa$$word
```

Source: Cisco ASA firewall

Figure 4-6 Unencrypted passwords in the show running-configuration command output

Figure 4-7 shows the same output after the service password-encryption command has been executed. The number "7" before the encrypted passwords indicates the level of encryption used.

```
line con 0
 password 7 03345A4F42182E5E4A
 login
line aux 0
 password 7 08114D0A4D0E0A0516
line vty 0 4
 password 7 08114D0A4D0E0A0516
```

Source: Cisco ASA firewall

Figure 4-7 Encrypted passwords in the show running-configuration command output

Banners

Banners are messages displayed to greet users who log on to a router. Banners provide information or warnings during logons, during privileged exec mode processes, or for an incoming asynchronous line connection. The most common banners display legal disclaimers and warnings at logon, and all banners should include a legal warning that clearly states the company's policy on unauthorized access.

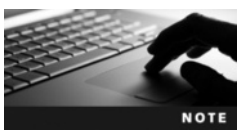

Legal disclaimers can cover complex issues, so legal counsel should be consulted to write or approve the disclaimers if your budget allows it. You need to address different jurisdictions because they might require varying notifications, particularly when transmissions are monitored.

A typical banner might look like this:

> WARNING! Authorized access only. Unauthorized access prohibited.
>
> This system is the property of [*Company Name*]. Disconnect IMMEDIATELY if you are not an authorized user!

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

A banner might be more explicit. For example, the following banner is modified from the U.S. Department of Energy's Computer Incident Advisory Capability (CIAC):

> THIS IS A PRIVATE COMPUTER SYSTEM. It is for authorized use only.
>
> Users (authorized or unauthorized) have no explicit or implicit expectation of privacy.
>
> Any or all uses of this system and all files on this system may be intercepted, monitored, recorded, copied, audited, inspected, and disclosed to authorized site and law enforcement personnel as well as authorized officials of other agencies, both domestic and foreign. By using this system, the user consents to such interception, monitoring, recording, copying, audition, inspection, and disclosure at the discretion of authorized site personnel.
>
> Unauthorized or improper use of this system may result in administrative disciplinary action and civil and criminal penalties. By continuing to use this system you indicate your awareness of and consent to these terms and conditions of use. LOG OFF IMMEDIATELY if you do not agree to the conditions stated in this warning.

A banner should never include wording that could give attackers information about your system or network, such as names, IP addresses, and software versions. Avoid displaying any information that an attacker might seek in network reconnaissance.

Remote Access with Secure Shell

A major security problem with routers, especially remote routers, is that network sniffers can intercept packets. The problem is compounded with connections to routers because they usually control the gateway access to your network. Protocols such as Telnet and FTP send everything in clear text, making them extremely vulnerable to packet sniffers. Standard remote shell programs, such as rlogin, rsh, and rcp, are just as susceptible to attack and should never be used to connect to a device outside your physically controlled network. To manage large network infrastructures, you can use **Secure Shell (SSH)**, a remote shell program, instead of Telnet or FTP.

SSH-2 is an update of SSH-1 and is more secure. Instead of being a single protocol like SSH-1, SSH-2 is a suite of protocols. In addition, SSH-1 and SSH-2 encrypt at different parts of packets. Another difference is that SSH-1 uses server and host keys to authenticate systems; SSH-2 uses only host keys. SSH-1 is not widely used now because SSH-2 has stronger security.

An alternative to the proprietary SSH is OpenSSH, which is said to be more secure because of the open nature of the code. OpenSSH includes several tools: secure copy (scp) as a replacement for rcp, secure ftp (sftp) as a replacement for standard FTP, and the SSH daemon (sshd). Several open-source variations are available for UNIX, Windows, and multiplatform networks.

Setting up and configuring SSH servers and clients are beyond the scope of this book because of the complexities of installing them on different operating systems; however, you can find more information at *www.ssh.com* and *www.openssh.com*.

SSH-1 Server was introduced on some Cisco devices with IOS 12.0.5.5, and the SSH client was introduced with IOS 12.1.3.T. Although only SSH-1 is used in the Cisco IOS, support for SSH-2 was added to some implementations beginning with IOS 12.1.(19)E. The current Cisco IOS offers an SSH server and client so that you can connect to your router's SSH server from an SSH client or connect your router's SSH client to another SSH server. If you

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

have routers that do not support SSH, you need to upgrade your router software. At Cisco's Web site, search for your router model and IOS version, and download the upgrade. You can find full instructions for installing as well.

Enabling SSH on the Router Before enabling SSH, you must configure the router with a hostname and domain name, and configure at least one interface with a static IP address. To determine whether your router supports SSH, issue the show flash command to verify your IOS version, and then compare it to the current list at Cisco's Web site or do an online search. After these preliminary steps, you can enable the SSH server on your router by using the following command to generate an RSA key pair:


```
Branch06(config)#crypto key generate rsa
```

This command enables SSH automatically; deleting the RSA key disables SSH automatically. Next, you choose a key size. Key sizes range from 360 to 2048, but be sure to use a key that is larger than the default size of 512 to ensure strong encryption. (A key size of 1024 should work for most applications.) To create new host keys, issue the following command again to overwrite the previous keys:

```
Branch06(config)#crypto key generate rsa
```

To delete all RSA keys and certificates and disable SSH, use the following command:

```
Branch06(config)#crypto key zeroize rsa
```

After the key pair has been generated on the router, SSH is enabled. However, a few more steps are necessary to fine-tune your SSH server. First, you should configure your authentication timeout interval—the time in seconds the server waits for a client to respond with a password. Enter the following command:

```
Branch06(config)#ip ssh time-out 60
```

The maximum (and default) setting is 120 seconds, but 60 seconds was configured in the preceding example. The next step is to configure the number of logon attempts allowed before the router drops the connection. Enter the following command:

```
Branch06(config)#ip ssh authentication-retries 3
```

The maximum is five attempts, and the default is three.

Create a user account with the following command syntax:

```
Branch06(config)#username [username] [priv] [priv level] [pass] [password]
```

The following list explains the parameters in this syntax:

- username—This is the command.
- [*username*]—This represents the logon name.
- [priv]—This optional command allows you to set the access level for the user.
- [*priv level*]—This field is used to set the access level. For example, use 15 to grant administrative privileges.
- [pass]—This indicates that the user's password follows.
- [*password*]—This is the user's new password.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

For example:

```
Branch06(config)#username dean priv 15 pass Pa$$word
```

Once the user has been created, you must allow logins on the VTY lines as follows:

```
Branch06(config)#line vty 0 4
Branch06(config-line)#login local
```

To connect to a router using SSH, the connecting systems need to have SSH client software installed. Several products can be used for this purpose. A popular choice is PuTTY, which works on UNIX and Windows systems. PuTTY still supports SSH-1 but now has support for SSH-2. When a PuTTY SSH client connects to the router for the first time, it is asked to approve caching the router's public encryption key, as shown in Figure 4-8.

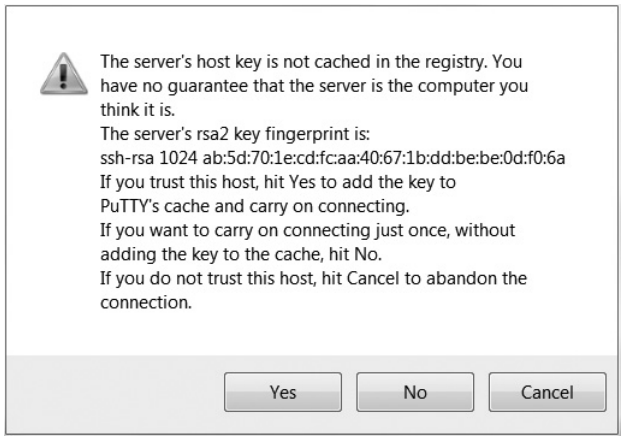

Source: PuTTY

Figure 4-8 PuTTY security alert

Figure 4-9 shows a packet capture of the connection between a PuTTY SSH client running on Windows 7 and a Cisco 2600 router running IOS version 12.4.

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	192.168.1.201	192.168.1.200	TCP	66	49855 > ssh [SYN] Seq=0 Win=8192 Len=0 MSS=1460 WS=4 SACK_PERM=1
2	0.004869	192.168.1.200	192.168.1.201	TCP	60	ssh > 49855 [SYN, ACK] Seq=0 Ack=1 Win=4128 Len=0 MSS=1460
3	0.005158	192.168.1.201	192.168.1.200	TCP	54	49855 > ssh [ACK] Seq=1 Ack=1 Win=64240 Len=0
4	0.010868	192.168.1.200	192.168.1.201	SSHv2	74	Server Protocol: SSH-1.99-Cisco-1.25
5	0.011150	192.168.1.201	192.168.1.200	SSHv2	82	Client Protocol: SSH-2.0-PuTTY_Release_0.62\r
6	0.011229	192.168.1.201	192.168.1.200	TCP	566	[TCP segment of a reassembled PDU]
7	0.011263	192.168.1.201	192.168.1.200	SSHv2	182	Client: Key Exchange Init
8	0.015512	192.168.1.200	192.168.1.201	SSHv2	334	Server: Key Exchange Init
9	0.021802	192.168.1.201	192.168.1.200	SSHv2	198	Client: Diffie-Hellman Key Exchange Init
10	0.330018	192.168.1.201	192.168.1.200	SSHv2	198	[TCP Retransmission] Client: Diffie-Hellman Key Exchange Init
11	0.377552	192.168.1.200	192.168.1.201	TCP	60	ssh > 49855 [ACK] Seq=301 Ack=813 Win=3316 Len=0
12	0.763002	192.168.1.200	192.168.1.201	SSHv2	502	Server: Diffie-Hellman Key Exchange Reply
13	0.765272	192.168.1.200	192.168.1.201	SSHv2	70	Server: New Keys
14	0.765345	192.168.1.201	192.168.1.200	TCP	54	49855 > ssh [ACK] Seq=813 Ack=765 Win=63476 Len=0
15	0.787297	192.168.1.201	192.168.1.200	SSHv2	70	Encrypted request packet len=16
16	0.787538	192.168.1.201	192.168.1.200	SSHv2	142	Encrypted request packet len=88
17	0.792022	192.168.1.200	192.168.1.201	SSHv2	106	Encrypted response packet len=52
18	1.000042	192.168.1.201	192.168.1.200	TCP	54	49855 > ssh [ACK] Seq=917 Ack=817 Win=63424 Len=0
19	3.224022	192.168.1.201	192.168.1.200	SSHv2	158	Encrypted request packet len=104
20	3.227413	192.168.1.200	192.168.1.201	SSHv2	106	Encrypted response packet len=52
21	3.423134	192.168.1.201	192.168.1.200	TCP	54	49855 > ssh [ACK] Seq=1021 Ack=869 Win=63372 Len=0

Source: Wireshark

Figure 4-9 Packet capture of an SSH connection

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

In packets 1, 2, and 3, the SSH client at 192.168.1.201 completes a TCP three-way handshake with the Cisco router at 192.168.1.200. In packets 4 and 5, the two systems identify themselves. In packets 7 through 14, the cryptographic keys are exchanged. Finally, in packets 15 through 21, the SSH session continues as encrypted router commands are sent to the router.

Verifying SSH Verifying that the router's SSH server is enabled is critical to security. In privileged exec mode, issue the show ip ssh command. The following example shows the command and the output it returns:


```
Branch06#show ip ssh
SSH Enabled - version 1.99
Authentication timeout: 60 seconds; Authentication retries: 3
```

If the SSH server is not enabled, you see this output:

```
SSH Disabled - version 1.99
Please create RSA keys to enable SSH.
```

You can also verify connections to the SSH server by using the show ssh command. If you encounter problems when checking your router, verify that you have specified a hostname and domain. Also, make sure you have generated an RSA key pair correctly by using the crypto key generate rsa command. If you have difficulty connecting to the router, you might not have configured the router's VTY lines by setting a password on the VTY interface. Remember that without a password, the router rejects any attempts to connect to the VTY interfaces.

You should also set a session timeout on VTY interfaces to reduce the risk of administrators leaving their computers unattended while logged on. Use the following exec-timeout command to set a timeout value of 10 minutes:

```
Branch06(config-line)#exec-timeout 10 0
```

Hardening a Router

A router is not much different from a standard desktop system, in that it has a processor, RAM, and an OS, and you can secure (harden) it as you would any server or workstation by removing unnecessary protocols and services. As a general rule, your router should support only required traffic. In other words, you should disable any unnecessary service or protocol.

The best place to start when determining what services and protocols to disable is your router security policy. This document should specify what traffic is allowed and whether traffic is incoming or outgoing. If a service or protocol is needed on one interface, check to see whether it can be disabled on other interfaces. If they are not needed, you should consider disabling services such as bootp (Bootstrap Protocol), which issues IP configuration to network clients, or CDP (Cisco Discovery Protocol), which shares information between Cisco devices. Also, remember to check the router vendor's Web site for new patches and security notices. You can find Cisco security bulletins at *http://tools.cisco.com/security/center/publicationListing*. The CERT Coordination Center (*www.cert.org*) is a helpful source of information about new security issues.

Logging is an important component of router security. As you learned in the logging section earlier in this chapter, centralizing logging makes it more difficult for attackers to remove evidence of their unauthorized presence. While they might be able to exploit a router, they would also have to compromise the logging host to cover their tracks. In addition, if the logging

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

host maintains its security logs on another server, the process of covering tracks becomes even more difficult.

Configuration management is the process of formally proposing, evaluating, approving, and implementing router configuration changes. This management helps avoid security vulnerabilities that result from casual configuration changes and allows systematic restoration to previous settings.

Patching routers has become somewhat controversial. Administrators once believed that attackers needed to know exactly what version of the Cisco IOS was running on a router to be able to execute code remotely. Because determining the IOS version is relatively difficult, conventional wisdom dictated that patching router software with updates created more risk of network outages than of preventing remote attacks. In the past few years, it has become clear that some methods of attack do not require knowledge of the IOS version. Administrators have discovered attacks caused by stack overflow errors and attacks on SNMP because ACLs were not configured properly. Careful ACL configuration and software patching are effective against such attacks.

As you have learned in this chapter, following security practices such as authentication, authorization, and accounting management, as well as logging, encryption, and access control lists, are standard procedures for well-managed routers.

Chapter Summary

- Routers direct the transportation of packets across networks. Because routers contain important information about a network's topology, you must consider their functionality and security carefully.
- Routers process OSI Network layer headers to determine the source and destination addresses for routing packets. Based on the information in the routing table, the router determines which of its multiple interfaces to use to move the packet closer to its destination.
- There are several ways to access a router for administrative purposes. The AUX port is used for dial-in serial connections. The CON port is used when the administrative computer is in the same physical location as the router. A rollover cable is used for a CON connection. Protocols like Telnet or the more secure SSH are used to access a router over a network—Ethernet, for example.
- Routing tables contain information about the network topology and are stored in the router's memory. This information can be acquired statically when an administrator enters the information manually, or it can be acquired dynamically when routers are configured to share information with each other automatically over the network.
- Static routing saves network bandwidth and gives administrators fine-grained control over small networks, but it is insufficient for larger networks. Dynamic routing protocols are automatic, and they can adjust to changing conditions on a network much faster and more efficiently than a network administrator can, but they also use network bandwidth as the routers communicate with each other.
- Various routing protocols such as RIP, OSPF, EIGRP, and IS-IS are used in LANs.
- To minimize the number of routes that must be maintained in routers (especially Internet routers), routes can be summarized through the process of supernetting, in which the subnet mask is adjusted to allow a single route to represent a number of subnets.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Access control lists are created to allow routers to perform packet filtering. Routers can filter packets based on source and destination IP addresses and ports, as well as protocols, TCP flags, IP options, and other factors. An access control list is applied to a router's interface to filter either inbound or outbound traffic. Only one ACL per interface, per direction, per protocol is permitted.
- Logging packet filtering and configuration activity is an important part of router and network security.
- Authentication, authorization, and accounting must be managed carefully to ensure router security and accessibility.
- Password security is not particularly strong on Cisco routers, which makes auditing all the more important. You can set passwords on the various access lines to the router and set passwords for privileged exec mode.
- Older router access methods such as Telnet are not secure because data is transferred in clear text. Today, the use of encrypted access methods such as SSH is common.
- Routers should be hardened in the same way as servers and other computers. Unnecessary services, accounts, and protocols should be removed.

Key Terms

access control entry (ACE) An individual rule in an ACL.

access control list (ACL) A group of rules that permit or deny passage of packets through a router based on parameters such as source and destination IP address, port address, protocol time to live, IP options, and TCP flags.

Address Resolution Protocol (ARP) A networking protocol that resolves IP addresses to MAC addresses.

antispoofing Measures to address digital attacks in which false source and destination addresses are used to bypass a router's packet filters.

ARP table A network system tool that lists the MAC and IP address resolutions of other devices on the network, making the resolution process more efficient.

authentication, authorization, and accounting (AAA) logging The process of verifying users' identities, establishing their level of access, and recording their activity with the use of a RADIUS or TACACS+ server.

auxiliary (AUX) port The port on a Cisco device that permits dial-in access to the device.

banner A message, usually a warning about appropriate use, presented to users of a digital system before authentication.

buffered logging The process of outputting logged data to a Cisco router's RAM.

Cisco Internetwork Operating System (IOS) The operating system used on Cisco devices such as routers, switches, and firewalls.

console (CON) port The port on a Cisco device that permits direct physical access from a nearby computer using the serial RS-232 protocol.

convergence A state in which all routers on a network have up-to-date routing tables.

default routes Manual entries in a router's routing table that specify how to handle packets that are not routable.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

distance-vector routing protocol A protocol that uses simple distance measurements such as hops to determine the best route for a packet.

dynamic routes Entries in a router's routing table that are generated automatically by communication through protocols with other routers on the network.

inverse mask The number in an access control list that specifies which part of an IP address is considered significant.

link-state routing protocol A routing protocol that uses complex calculations to determine the best route for a packet.

logging The process of recording activity on a computing device.

metrics Cost values that help routers assess the desirability of a link.

named ACLs An advancement over extended and standard ACLs that allow filtering packets based on IP options and TCP flags, and that permit insertion and modification of individual ACEs.

Remote Authentication Dial-In User Service (RADIUS) An authentication, authorization, and accounting service protocol.

rollover cable A cable used to connect to a Cisco CON port from a computer's COM port via the serial RS-232 protocol.

route summarization The process of shortening the network identifier in a subnet mask to allow one network route to represent multiple network routes. (Also called *supernetting*.)

routed protocol A networking protocol used to transport data packets.

routing The process of transporting packets of information across a network from the source node to the destination node.

routing protocol A protocol used by routers to share their routing tables with each other.

routing table The information about accessible networks that is maintained by a router.

Secure Shell (SSH) A protocol used to secure remote sessions by using asymmetric encryption.

Simple Network Management Protocol (SNMP) A protocol used to monitor the status of network devices.

SNMP trap logging A specific parameter of network activity to monitor.

standard ACLs Access control lists that are limited in filtering parameters to a source IP address or network.

static routes Manual entries in a router's routing table by an administrator.

stub network A network segment on a larger network's edge; a dead-end segment.

stub router A router that connects a stub network to the larger network.

supernetting See *route summarization*.

syslog server A log management service.

Terminal Access Controller Access Control System Plus (TACACS+) An authentication, authorization, and accounting service protocol.

variable length subnet masking (VLSM) The process of adjusting the subnet mask within a single network to allocate IP addresses more efficiently.

virtual terminal (VTY) Virtual session access points for simultaneous access to a Cisco device.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Review Questions

1. Consider the following statements:

   A: Dynamic routing protocols decrease network security.

   B: Static routing protocols conserve network bandwidth.

   a. Statement A is true, and statement B is false.
   b. Statement A is false, and statement B is true.
   c. Statement A is true, and statement B is true.
   d. Statement A is false, and statement B is false.

2. A packet has a destination Data Link layer address of FF-FF-FF-FF-FF-FF.
   a. Any router that receives this packet forwards it to the default route interface.
   b. Any nonrouting device drops this packet.
   c. Network hosts send this packet to their default gateway.
   d. Network hosts process this packet to see if it is of interest.

3. Which of the following addresses are valid IPv6 addresses? (Choose all that apply.)
   a. fe80::a02a:64b0:27a9:c73b
   b. fe80:cd5c::f40f:9eea:7580
   c. fe80::bef2:6cbe::678:1879
   d. fe80:c693:bef2:6cbe:678:1879

4. Which of the following features are typical of dynamic routing protocols? (Choose all that apply.)
   a. increased administrative effort
   b. faster convergence
   c. decreased processor load
   d. increased network traffic

5. Consider the following statements:

   A: Supernetting reduces the number of routes in routing tables.

   B: Variable length subnet masks allow routers to determine whether an address is public or private.

   a. Statement A is true, and statement B is false.
   b. Statement A is false, and statement B is true.
   c. Statement A is true, and statement B is true.
   d. Statement A is false, and statement B is false.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

6. All Cisco access control lists require _____.
   a. at least one permit statement
   b. at least one deny statement
   c. either a source or destination IP address in each access control entry
   d. a deny all statement as the last access control entry

7. Which of the following parameters can you find in a standard access control list? (Choose all that apply.)
   a. source IP address
   b. destination IP address
   c. dource port number
   d. destination MAC address

8. Which of the following IP addresses would be filtered by an access control entry that contained the IP address and inverse mask specification of 12.96.115.77 0.15.255.255? (Choose all that apply.)
   a. 21.48.200.9
   b. 12.104.146.190
   c. 12.111.115.77
   d. 12.113.84.3

9. Which of the following packet elements can be filtered by a named ACL? (Choose all that apply.)
   a. TCP flags
   b. noninitial fragments
   c. destination ports
   d. source IP addresses

10. Cisco encryption type 7 is the strongest encryption method available on the router. True or False?

11. To configure a router using a telephone connection, you would connect to the _____ port.
   a. CON
   b. AUX
   c. TTY
   d. VTY

12. To enable SSH on a Cisco router, you should _____.
   a. run the command `ssh enable rsa`
   b. run the command `crypto key zeroize rsa`
   c. create RSA keys
   d. install SSHv3

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

13. Which of the following are routing protocols? (Choose all that apply.)
    a. EIGRP
    b. ARP
    c. MAC
    d. RIP

14. Protocols that guide a router's decisions on packet forwarding based on the current condition of the network are _____ .
    a. link-state protocols
    b. static protocols
    c. the most secure routing protocols
    d. routed protocols

15. Consider the following statements:

    A: An ACL can be applied only to one router interface.

    B: An ACL can apply only to one routed protocol.

    C: An ACL can filter packets traveling in only one direction.
    a. Statements A and C are correct.
    b. Statements A and C are incorrect.
    c. Statement A is correct, and statement B is false.
    d. Statement A is false, and statement B is correct.

Hands-On Projects

Hands-On Project 4-1: Downloading and Installing a Router Simulation Program

Time Required: 10 minutes

Objective: Download and install a router simulation program.

Description: You download and install a free router simulation program in this project. Most router configuration is done from a command-line interface, and the program you install in this project gives you practice in working at the Cisco IOS command line.

1. Log on to either the Windows 7 or Windows Server 2008 computer. You can perform all the projects in this chapter on either system.
2. Open a Web browser and go to *http://download.cnet.com/Toggit-Router-Simulator-and-Lab/3000-2051_4-10110264.html*.

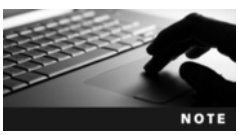

Because the availability of freeware programs can change, you might have to find a different router simulator for these projects. You can do these projects no matter what program you use.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

3. Click the green **Download Now** link. Click **Save**. When the download is complete, click **Run**. If you are using Internet Explorer 9, click **Save**, click the **View downloads** button while the file is downloading, and then click **Run** when the download is finished. If a User Account Control window appears, enter your username and password, and then click **Yes**.
4. With the CNET Download.com Installer window, you might also see a Special Offer window. If so, click **I do not accept**, and click **Next Step**. When the download is complete, click **Open**. Double-click the **ToggitSim** folder, and double-click **setup**. If a User Account Control window opens, enter the username and password of the domain administrator, and then click **Yes**.
5. In the Toggit Router Simulator Setup window, click **OK**. Note the directory in which Toggit will be installed; by default, this directory is C:\Program Files(x86)\Toggit\RouterSim. Click the icon that shows the computer and box to start the installation. In the Choose Program Group window, click **Continue**. When the setup program has completed successfully, click **OK**.
6. You might be able to open the program by clicking **Start**, clicking **All Programs**, clicking the **Toggit** folder, and clicking **Router Simulator**. In some cases, the program does not create a Start menu link. If so, browse to the installation folder that you noted in Step 5, right-click **ccnasim**, and click **Pin to Start Menu**. Start the program by clicking **Start** and clicking **ccnasim**. Click **Start**. The opening screen should look like Figure 4-10.

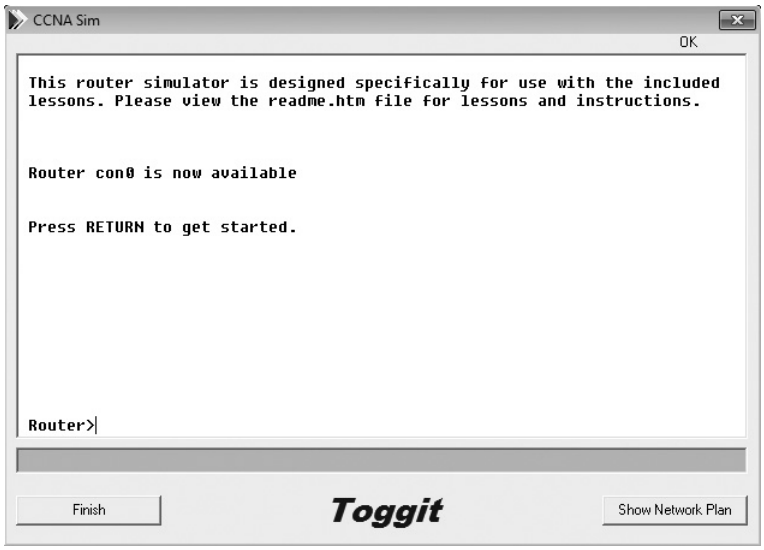

Source: CNET/Toggit

Figure 4-10 A typical router simulation program

7. Leave the router simulation program running for the next project.

Hands-On Project 4-2: Working at the IOS Command-Line Interface

Time Required: 10 minutes

Objective: Learn basic router command-line modes and navigate among them.

Description: In this project, you learn how to log on to a router and move among different EXEC modes. You learn basic commands for viewing configuration information and how

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

to navigate among configuration modes to configure, maintain, and troubleshoot Cisco routers.

1. Return to your router simulation program from Hands-On Project 4-1. You should see the > prompt, which is called the user mode prompt. From this prompt, you can view certain settings and enter privileged mode, where you can modify the router's configuration. Type **?** to see the commands available in this router's user mode. Not all router simulators provide this feature, but it is available on all Cisco routers.
2. Type **enable** and press **Enter**. If your router simulator asks for a password, you can press **Enter** because you have not configured one. You should see a # prompt to indicate that you are in privileged mode, where you can change the router's name, program banners, set passwords, and enter other configuration modes.
3. Type **configure terminal** and press **Enter**. The prompt changes to Router(config)#, which is global configuration mode. Type **hostname Branch06** and press **Enter**. You have named the router, so the prompt changes to Branch06(config)#.
4. Type **interface ethernet 0** and press **Enter**. The prompt changes to Branch06(config-if)#, which is interface configuration mode for the Ethernet 0 interface. Type **exit** and press **Enter**. You are now in interface configuration mode for the serial 0 interface. Press **Ctrl+Z** to return to privileged mode.
5. Type **show running-config** and press **Enter**. Examine the output, which shows information about the IOS version, hostname, services, line, and interfaces. Return to global configuration mode, as indicated by the Branch06(config)# prompt, type **line vty 0 4** and press **Enter**. The prompt changes to Branch06(config-line)#.
6. Press **Ctrl+Z** to return to privileged mode, and then type **exit** and press **Enter** to return to user mode.
7. Leave your router simulation program running for the next project.

Hands-On Project 4-3: Assigning IP Addresses to Interfaces

Time Required: 10 minutes

Objective: Configure IP addresses for Ethernet and serial router interfaces.

Description: In this project, you configure the Ethernet 0 and serial 0 router interfaces with IP addresses. You then use the show commands to see how the interfaces change as you work with them.

1. Enter privileged mode (the Branch06# prompt). Type **show interface ethernet 0** and press **Enter**. Notice that the Ethernet 0 interface is up but the line protocol is down. Your result should look similar to Figure 4-11.

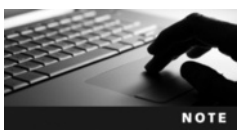

In the show interface output, the Ethernet interface refers to the router's physical Ethernet interface, and the line protocol refers to the line carrying the signal. If the line is usable, the field states that the line protocol is up, meaning that keep-alives are successful. The meanings of "up" and "down" are straightforward: "Up" means that the protocol is working, and "down" means it is not working or that the administrator has taken it down, in which case the output is "administratively down."

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

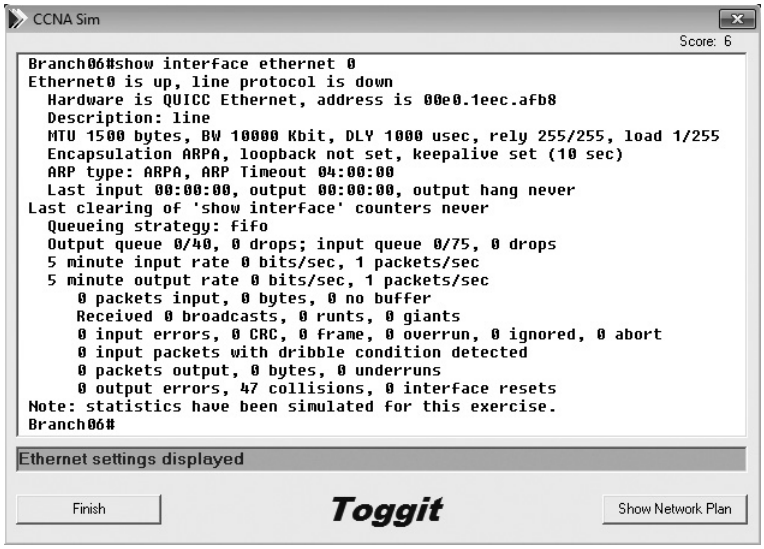

Source: CNET/Toggit

Figure 4-11 Sample output from the show interface ethernet 0 command

2. Enter global configuration mode by typing **configure terminal** and pressing **Enter**. Type **interface ethernet 0** and press **Enter** to switch to Ethernet 0 interface configuration mode, prompt Branch06(config-if)#. Type **ip address 172.10.0.1 255.255.0.0** and press **Enter**. Type **no shutdown** and press **Enter** to enable the interface.

Cisco routers allow you to truncate many commands. For example, a real router would accept "en" instead of "enable" or "config t" in place of "configure terminal." Some simulation programs also accept shortened commands, but because many do not, complete commands are used in these projects.

3. Exit to privileged mode, type **show interface ethernet 0** and press **Enter**. The output now shows that Ethernet 0 is up and the line protocol is up. The IP address you just programmed is displayed along with statistics for the interface.
4. Return to interface configuration mode for Ethernet 0, type **shutdown** and press **Enter**. Exit back to privileged mode and check the show interface output for Ethernet 0 again. Now Ethernet 0 is up, but the line protocol is down.
5. Enter interface configuration mode for the Serial 0 interface. Type **ip address 210.16.54.1 255.255.255.0** and press **Enter** to assign the IP address to the interface. Serial interfaces are a little different from Ethernet interfaces. In some cases you must also set the clock rate to synchronize the transmissions. Type **clock rate 56000** and press **Enter**. Enable the interface by typing **no shutdown** and pressing **Enter**, and then press **Ctrl+Z** to exit back to privileged mode.
6. Type **show running-config** and press **Enter**. Note that Ethernet 0 and Serial 0 now show IP addresses.
7. Leave your router simulation program running for the next project.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Project 4-4: Configuring Static Routes

Time Required: 10 minutes

Objective: Configure a static route.

Description: Use the router simulation program to enter a static route in the routing table, and then view the routing table.

1. From privileged mode, type **show ip route** and press **Enter**. If you completed Hands-On Projects 4-2 and 4-3, you should see a "C" before the IP address for the Ethernet and Serial interfaces, indicating that these networks are directly connected to the router.
2. Enter global configuration mode. Type **ip route 192.168.10.0 255.255.255.0 210.16.54.1** and press **Enter**. This command tells the router that the network 192.168.10.0, which uses the subnet mask of 255.255.255.0, can be reached by sending the packet its own interface at 210.16.54.1. Enter another network destination by typing the following command: **ip route 192.168.50.0 255.255.255.0 210.16.54.1**. Press **Enter**. This command indicates that two different networks are accessible through Branch06 router's 210.16.54.1 interface. Figure 4-12 shows a network with these routes.

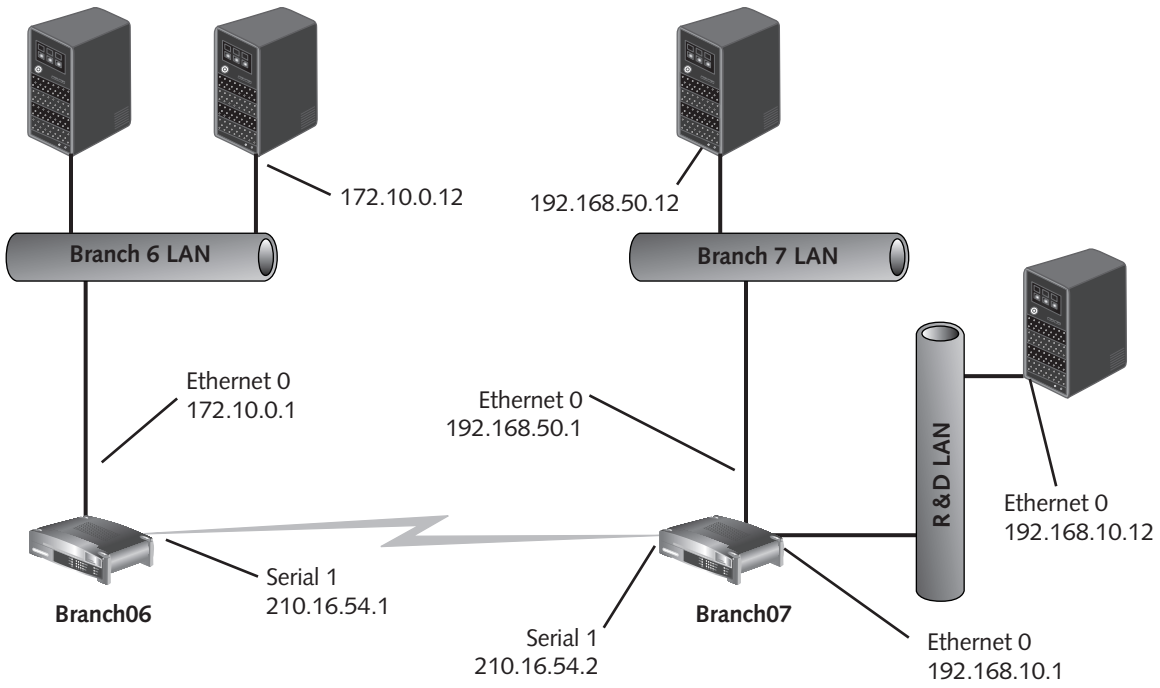

© Cengage Learning 2014

Figure 4-12 Routed network

3. Exit to privileged mode. Type **copy running-config startup-config** and press **Enter**. This command saves your configurations so far to nonvolatile RAM. If the router were to

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

lose power, it would "remember" all your previous configurations when it was turned on again. Before you typed the preceding command, your configurations were stored only in volatile RAM and would have been lost if the router had lost power.

4. Type **show ip route** and press **Enter**. Two additional routes should be displayed. The "S" indicates that the routes have been statically entered. (See Figure 4-13.)

```
Branch06# show ip route
C 172.10.0.1 is directly connected, Ethernet0
C 210.16.54.0 is directly connected, Serial0
S 192.168.50.0 [1/0] via 210.16.54.1
S 192.168.10.0 [1/0] via 210.16.54.1
```

Source: Cisco ASA firewall

Figure 4-13 Output from the show ip route command

5. Click **Finish** to close the router simulation program.

Hands-On Project 4-5: Exploring Windows IPv6 Implementation with Netsh

Time Required: 15 minutes

Objective: Explore Windows IPv6 using the netsh command.

Description: The Windows netsh command is very powerful. You can use it to do everything from setting specific configurations for networking protocols to rebuilding your TCP/IP implementation. In this project, you use the netsh command to examine the IPv6 configuration on your Windows system.

1. Log on to Windows 7 or to Windows Server 2008 with administrative accounts.
2. Click **Start**. In the Search programs and files box, type **cmd** and then press **Enter**.
3. Type **netsh interface ipv6 show address** and press **Enter**. The result should look similar to Figure 4-14. In this project, your results differ to varying degrees depending on your IP configuration and your network configuration.

```
C:\Users\Administrator>netsh interface ipv6 show address

Interface 1: Loopback Pseudo-Interface 1

Addr Type  DAD State   Valid Life Pref. Life Address
---------  ----------- ---------- ---------- ------------------------
Other      Preferred     infinite   infinite ::1

Interface 12: isatap.{64CC01BF-4F74-46FF-8089-A64ED4BC4F68}

Addr Type  DAD State   Valid Life Pref. Life Address
---------  ----------- ---------- ---------- ------------------------
Other      Deprecated    infinite   infinite fe80::5efe:192.168.1.110%12

Interface 11: Local Area Connection

Addr Type  DAD State   Valid Life Pref. Life Address
---------  ----------- ---------- ---------- ------------------------
Other      Preferred     infinite   infinite fe80::8461:45c7:a318:8316%11

Interface 13: Local Area Connection* 11

Addr Type  DAD State   Valid Life Pref. Life Address
---------  ----------- ---------- ---------- ------------------------
Other      Deprecated    infinite   infinite fe80::ffff:ffff:fffe%13
```

Source: Microsoft Windows 7

Figure 4-14 Output of the netsh interface ipv6 show address command

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

4. Type **ipconfig /all** and press **Enter**. Compare the two results. You may find that some of your IPv6 addresses are associated with virtual network interfaces. In Figure 4-15, only the interface with the address fe80::8461:45c7:a318:8316 is an actual network interface. You are likely to have addresses that begin with fe80. What types of addresses are these and how did they get there?
5. Type **netsh interface ipv6 show route** and press **Enter**. Your result should look similar to Figure 4-15.


```
C:\Users\Administrator>netsh interface ipv6 show route

Publish  Type      Met  Prefix                    Idx  Gateway/Interface Name
-------  --------  ---  ------------------------  ---  ------------------------
No       Manual    8    ::/0                       13  Local Area Connection* 11
No       Manual    256  ::1/128                     1  Loopback Pseudo-Interface 1
No       Manual    8    2001::/32                  13  Local Area Connection* 11
No       Manual    256  fe80::/64                  11  Local Area Connection
No       Manual    256  fe80::/64                  13  Local Area Connection* 11
No       Manual    256  fe80::5efe:192.168.1.110/128   12  isatap.{64CC01BF-4F74-46F
No       Manual    256  fe80::ffff:ffff:fffe/128   13  Local Area Connection* 11
No       Manual    256  fe80::8461:45c7:a318:8316/128   11  Local Area Connection
No       Manual    256  ff00::/8                    1  Loopback Pseudo-Interface 1
No       Manual    256  ff00::/8                   13  Local Area Connection* 11
No       Manual    256  ff00::/8                   11  Local Area Connection
```

Source: Microsoft Windows 7

Figure 4-15 Output of the netsh interface ipv6 show route command

Notice the listings for the interfaces. The /128 indicates that the entire address is identified, as opposed to subnet prefixes, which are indicated by the /8, /32, and /64. Notice also the listing of 2001::/32. What is the purpose of this route?

6. Type **netsh interface ipv6 show interfaces** and press **Enter**. Your result should look similar to Figure 4-16.

```
C:\Users\Administrator>netsh interface ipv6 show interfaces

Idx     Met         MTU          State                Name
---  ----------  ----------  ------------  ---------------------------
  1           50  4294967295  connected     Loopback Pseudo-Interface 1
 12           50        1280  disconnected  isatap.{64CC01BF-4F74-46FF-8089
 11           10        1500  connected     Local Area Connection
 13           50        1280  disconnected  Local Area Connection* 11
```

Source: Microsoft Windows 7

Figure 4-16 Output of the netsh interface ipv6 show interfaces command

As you can see in the figure, of the local area connections, only the one with the Index number of 11 is actually connected to the network. That explains why the most significant results in the previous command output contain the index number 11. The results also show that the system cannot have a "true" routing table because a system must have more than one interface connected to make routing decisions. The routing information is provided to enable IP on the system and determine whether packets should be sent to the network or to the gateway interface.

7. Type **netsh interface ipv6 show teredo** and press **Enter**. Your result should look similar to Figure 4-17. Teredo is a tunneling protocol supported by Microsoft that allows hosts on a non-native IPv6 network to communicate with the IPv6 Internet. The server name listing, teredo.ipv6.microsoft.com, is the system that connects the IPv4 and IPv6 Internet.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

```
C:\Users\Administrator>netsh interface ipv6 show teredo
Teredo Parameters
---------------------------------------------
Type                    : client
Server Name             : teredo.ipv6.microsoft.com.
Client Refresh Interval : 30 seconds
Client Port             : unspecified
State                   : offline
Error                   : client is in a managed network
```

Source: Microsoft Windows 7

Figure 4-17 Output of the netsh interface ipv6 show teredo command

8. Log off your system.

Case Projects

Case Project 4-1: Tunneling IPv6

You have been assigned to report to your network administrators on the use of Teredo. Prepare a two-page memo that outlines why Teredo was developed, how it is implemented in Windows operating systems, for which types of networks it is appropriate, and how long it should be implemented.

Case Project 4-2: Creating ACLs

You have been assigned to create access control lists to filter specific traffic on a Cisco router. Provide the commands needed to filter the appropriate traffic in each of the following ACLs.

ACL 1

- Allow Telnet connections to the 192.168.1.0 network from host 10.3.4.7.
- Allow established connections from network 172.16.0.0 to anywhere.
- Permit all other access.

ACL 2

- Prevent Telnet connections from network 192.168.1.0 to network 172.16.0.0.
- Prevent reserved addresses from accessing any network.
- Deny spoofing from the broadcast address.
- Permit all other access.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 5

Cryptography

After reading this chapter and completing the exercises, you will be able to:

- Explain components of cryptographic protocols
- Explain common cryptography standards
- Describe modern cryptanalysis methods

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

This chapter introduces you to the main concepts of cryptography. You learn about the components of cryptographic protocols, including cryptographic primitives, encryption algorithms, hashing algorithms, and digital signatures. They are the building blocks of modern protocols, and each component works with others to achieve the main goals of cryptography: confidentiality, message integrity, authentication, and nonrepudiation. Key management is a major challenge in cryptographic applications, so you also learn about private and public key exchange methods.

Next, this chapter reviews common implementations of cryptographic standards in information technology, including Data Encryption Standard and Advanced Encryption Standard. You learn about cryptographic applications used in wireless communication and on the Internet, and the encryption components of Internet Protocol Security (IPsec). Finally, you learn about modern cryptanalysis, including common attacks against encryption algorithms and cryptographic protocols, and ways to defend cryptographic systems against these threats.

Components of Cryptographic Protocols

Cryptography is the process of converting readable text, programs, and graphics into data that cannot be easily read or executed by unauthorized users. In other words, cryptography is the process of converting **plaintext** into **ciphertext** by using an encoding function such as an encryption algorithm or a secured list of substitution characters. The encrypted file can then be transmitted in many forms, including written messages and electronic data. Of course, the cryptographic process also includes providing a way for authorized users to access the plaintext. Cryptography helps achieve four critical goals of information security:

- *Confidentiality of information*—Only authorized users can access data.
- *Integrity of data*—Data has not been modified.
- *Authentication*—Users are who they claim to be.
- *Nonrepudiation*—Neither party can plausibly deny its participation in message exchanges. This is an important quality for business and legal transactions.

Cryptographic systems used in computer security contain common components, called cryptographic primitives, that are combined into cryptographic protocols and standards. **Cryptographic primitives** are modular mathematical functions that include encryption algorithms, hashing functions, pseudorandom number generators, and basic logical functions.

The more detailed mathematical aspects of cryptographic design are beyond the scope of this book. Learning the basics of how these protocols are structured can help you determine what types of protocols best serve your needs and understand their vulnerabilities and deployment requirements.

Cryptographic Primitives

Used alone, cryptographic primitives cannot provide data integrity, confidentiality, nonrepudiation, and authentication. A primitive can accomplish only one of these goals.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Each primitive is designed to perform a specific task reliably, such as generating a digital signature for a set of data. To provide adequate security, primitives must be used with other primitives.

In some ways, a cryptographic primitive is like a programming language. Software engineers do not create a new programming language for every new program. They use existing, proven languages in a modular fashion to create a software application that performs a defined set of functions. The same principle applies to cryptographic primitives. A cryptographic system designer does not create new primitives because they are complicated and prone to errors, even for experts.

Primitives are not usually the source of security failures in cryptographic protocols. Security flaws result from mistakes in designing the protocol, such as an overall poor design, poorly chosen primitive combinations, or bugs introduced during the design process.

Each primitive in a cryptographic system handles one aspect of securing data. For example, the encryption algorithm performs encoding, thus providing message confidentiality but not message integrity or authentication. By combining the encryption algorithm with a hashing function, however, you can meet the requirement of message integrity. You can also add a digital signature for message authentication.

In the following sections, you examine cryptographic primitives that are common to modern cryptographic systems. From the most basic logical functions to more intricate encryption algorithms and key generation routines, you learn the fundamentals of how primitives work on their own and when combined with others.

Exclusive OR Function The **exclusive OR (XOR) function** is used in cryptography as a linear mixing function to combine values. For example, the output of other primitive ciphers can be combined with an XOR function to produce a pseudorandom value on which another cipher performs additional operations. An XOR function is based on binary bit logic and results in a logical value of true if only one of the operands has a value of true. So, for example, if x and y are the same (both true or both false), the XOR output is 0 (false). If x and y are different, the XOR output is 1 (true). The truth table in Figure 5-1 shows inputs of x and y and the result of the XOR function in the right column.

x	y	xX0Ry
0	0	0
0	1	1
1	0	1
1	1	0

© Cengage Learning 2014

Figure 5-1 An XOR truth table

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The XOR function is useful as a cryptographic primitive because of its reversible property, as shown in the following example:

```
p XOR k = c
c XOR k = p
```

In this equation, *p* represents plaintext, *k* is the key, and *c* is ciphertext. The *c* resulting from `p XOR k` is reversible in the `c XOR k` statement.

Permutation Functions Bit-shuffling **permutation functions**, often used in symmetric algorithms, reorder sets of objects randomly—for example, by rearranging input bits, such as the binary input 010 into 001. A good analogy for a random permutation is shuffling a deck of cards, which ideally causes the deck to be dealt in a completely random order. Some variations of permutation functions are suitable for cryptographic use, but many are not. One cryptographically notable variation is an expansion permutation, in which certain bits are used more than once. For example, the input 010 is rearranged and expanded into 0101.

Substitution Box Functions A **substitution box (S-box) function** transforms a number of input bits into a number of output bits and produces a lookup table that can be fixed or dynamic, depending on the cipher. It is a basic component of symmetric key algorithms. The purpose of this function, as with permutation functions, is disguising the relationship of ciphertext to cleartext. An S-box function is usually described as *n* input bits × *m* output bits, so a 6×4 S-box means that 6 input bits are transformed into 4 output bits.

Feistel Networks Created in 1973, a **Feistel network** (shown in Figure 5-2) is a symmetric block cipher that is the basis of several symmetric encryption algorithms. A Feistel network's purpose is to obscure the relationship between ciphertext and keys (a shortcoming of symmetric algorithms). It does this by combining multiple rounds of repeated operations, such as processing cleartext input with XOR functions. A key schedule is used to produce different keys for each round. The advantage of a Feistel network is that its encryption and decryption operations are similar or even identical, which reduces the size of its code and the resources needed to use it. Feistel encryption works because the key schedule can be reversed, using keys in exact reverse order.

Pseudorandom Number Generators **Pseudorandom number generators (PRNGs)** are essential components of a cryptographic algorithm. A PRNG is an algorithm for generating sequences of numbers that approximate random values. To be considered a cryptographically secure pseudorandom number generator (CSPRNG), a function must meet certain design principles and be resistant to known attacks. Many cryptographic functions require random values that serve as seeds for further computation:

- *Nonces*—A nonce is a number or bit string (usually random) that prevents generation of the same ciphertext during subsequent encryptions of a message. Using nonces strengthens encryption and makes it more resistant to being broken.
- *One-way functions*—A one-way function is easy to compute but difficult and time consuming to reverse. One-way functions include integer factorization, discrete logarithms, and the Rabin function. They are considered one-way because no efficient

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

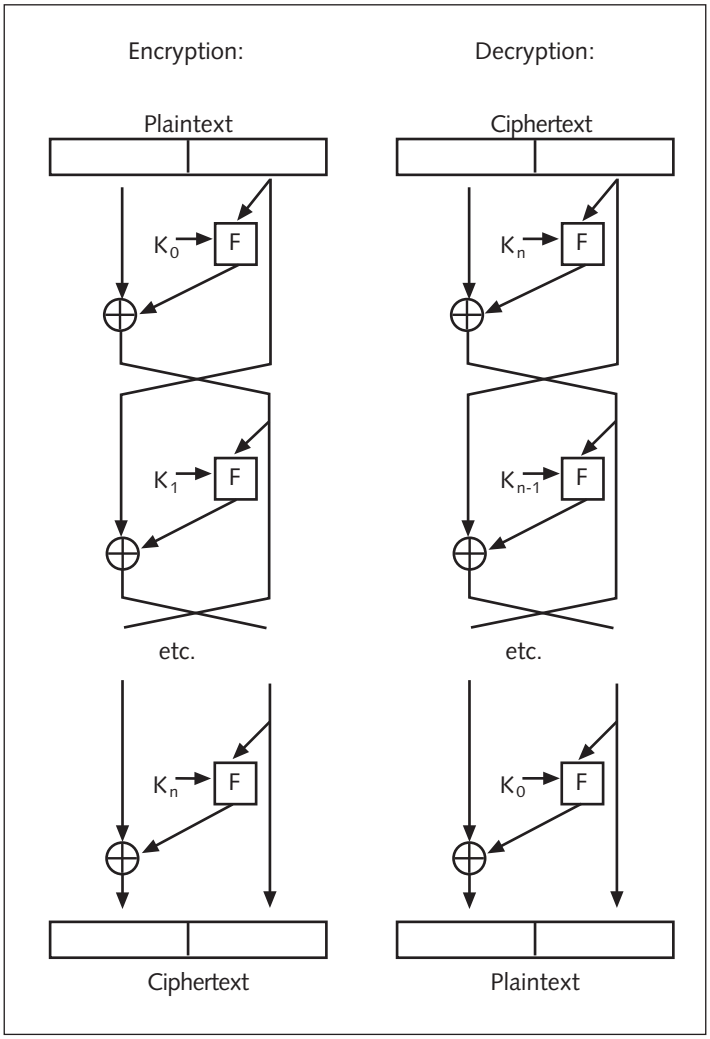

© Cengage Learning 2014

Figure 5-2 A Feistel network

inverting algorithm has been discovered, but new methods could prove these functions reversible.

- *Salts*—A salt consists of random bits used as input for key derivation functions and to pad values (hide the true contents). Initialization vectors, passwords, and passphrases are sometimes used as salts.
- *Key derivation (generation)*—A key derivation function generates secret keys from a secret value (usually a randomly generated value) and another piece of information such as a password.

PRNGs can be hardware or software based. In general, hardware PRNGs offer sequences that are closer to being truly random, but the hardware can be expensive and cumbersome

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

and is vulnerable to the same threats as any hardware security device. Several algorithms can produce random number sequences of sufficient quality for cryptographic purposes, making hardware PRNGs unjustifiably expensive in most cases.

PRNG values are not completely random because they are generated by a structured method based on fairly small initialization values called the PRNG's state. The state is generally measured in bits, and the initial state's size determines the maximum sequence length before it repeats; this is known as the PRNG's period. In modern computing, it is possible to provide a long enough state to ensure that the PRNG's period is longer than any computer could compute in a reasonable amount of time—say, 100 million years. If one bit is produced every picosecond, the PRNG's period with a 64-bit state is more than 140 million years.

PRNGs still have some problems, even if they are considered cryptographically secure. One problem is that by providing a starting point for random number generation, the resulting value is not truly random. Another problem is that the PRNG always produces the same values when initialized with a particular state. To be secure, mechanisms are necessary to prevent a PRNG from reusing an initializing state.

CSPRNGs produce values that are random enough to suit the intended use. In practice, there is no certain way to determine the output of a CSPRNG from truly random numbers without knowing the algorithm used and its initialization state. Because distinguishing between truly random and generated pseudorandom numbers is not considered feasible, most algorithms and protocols that use CSPRNGs are considered secure.

Hashing Functions One method of verifying message integrity is by using **hashing functions**, which generate a hash value or message digest from input. (Hashing algorithms, as you will learn later in this chapter, define the instructions for running hashing functions.) A hash value is a fixed-size string representing the original input's contents. If the input changes in any way, even by adding a period at the end of a sentence, the resulting output has a different hash value.

For messages sent over the Internet, verifying message integrity and authenticating the source are critical. Source authentication verifies the sender's identity and prevents messages from fraudulent or spoofed sources from being accepted. Using a hashing function to verify message integrity is as simple as comparing the message digest the sender calculates with the message digest the receiver calculates. If the values are the same, the sender's message has not been altered during transmission.

Hashing functions are also used for error detection, as with Cyclic Redundancy Check (CRC), a commonly used method of verifying that a message was not altered by a transmission error, such as interference on the transmission medium.

Encryption Algorithms

To see how cryptographic primitives are combined to provide confidentiality, integrity, authentication, and nonrepudiation, understanding the basic mathematical concepts is helpful. An algorithm is a precisely defined set of instructions for carrying out a task. Computer algorithms, for example, provide exact instructions for which operations to carry out, which criteria change operations, how many times to perform an operation (called looping), and when to stop.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

In algorithms, a strict order of operations is essential; in computer programs, this strict order of operations is called control flow. If a program processes the same input for the same purpose, but the instructions vary, the results would never match, nor would they be repeatable, reversible, and predictable. Output values change, of course, if the input varies, but the instructions for processing input must stay the same. For example, the instruction "Find the largest item in the list" generates different values depending on the list's contents, but must always process the list exactly the same way to achieve consistency.

An **encryption algorithm** is a set of precise instructions that provides an encoding function for a cryptographic system or generates output for use in additional operations. (Remember that some algorithms also double as CSPRNGs, such as stream and block ciphers.) In the simplest applications, an encryption algorithm is a mathematical formula that works with a key to generate ciphertext from cleartext input. Encryption algorithms also combine with other primitives that perform integrity checking or authentication. A hashing function, for example, can be used to check data integrity or to generate pseudorandom numbers that the encryption algorithm can use in encoding iterations. This process is repeatable and reversible.

Key Size in Encryption Algorithms An encryption algorithm's strength is often tied to its key length. The longer the key, the harder it is to break the encryption. Longer keys offer more protection against brute-force attacks, in which every possible key is tried to decrypt a message. As computers have increased in processing power, however, attackers have been able to carry out brute-force attacks more quickly and break encryption keys. Therefore, to keep up with advances in processors, encryption key sizes had to increase. As an illustration of the power of extending key size, consider a five-character password created from only lowercase letters. Such a password would offer almost 12 million possible combinations (26^5). However, extending the password to 10 characters would increase the number of possible combinations to more than 141 trillion (26^{10}).

Types of Encryption Algorithms The two major types of encryption algorithms are block ciphers and stream ciphers. A **block cipher** encrypts groups of text at a time. For example, a block cipher encrypts the whole word *cat* instead of encrypting each letter. A **stream cipher** encrypts cleartext one bit at a time to produce a stream of encrypted ciphertext, so the letters *c*, *a*, and *t* in *cat* are encrypted separately.

Block ciphers and stream ciphers use keys differently, so they fall into two categories: symmetric and asymmetric. **Symmetric algorithms** use the same key to encrypt and decrypt a message. They are considered the workhorses of the encryption world and are a faster, more efficient method of encrypting data because they require fewer computing resources. Most encryption protocols in daily use are based on symmetric algorithms, and the majority of symmetric algorithms are block ciphers.

By contrast, **asymmetric algorithms** use a specially generated key pair. One key encrypts cleartext into ciphertext, and the other key decrypts ciphertext into cleartext. Either of the generated pair can be used to encrypt, but the other key must be used to decrypt. Asymmetric encryption and decryption are about 10,000 times slower than symmetric encryption.

Blowfish Blowfish is a 64-bit block cipher composed of a 16-round Feistel network and key-dependent S-box functions. This unpatented cipher used worldwide has a variable key

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

size from 32 to 448 bits. The default key size is 128 bits. Blowfish is fast in encryption and decryption operations, but its 64-bit block size is now considered too short and makes Blowfish vulnerable to some attacks.

Notable for its public license status and excellent performance, Blowfish is still a widely used cipher, and is very fast except when changing keys. Its high memory requirements (4 kilobytes) and slow key derivation functions are not a problem for most desktop and laptop computers (even older ones), but these characteristics make it unsuitable for smart cards or similar limited-resource systems.

Twofish Twofish, the successor to Blowfish, is a 128-bit symmetric block cipher composed of a 16-round Feistel network and key-dependent S-box functions. Twofish also has a complicated key schedule and a variable key size of 128, 192, or 256 bits. Like Blowfish, it is publicly licensed, but it has not been used as much as its predecessor. Although some theoretical work on cryptanalysis methods against Twofish has been published, it has not been broken.

Rivest Cipher Family Rivest Cipher 4 (RC4) is a popular stream cipher in Web browsers that use Secure Sockets Layer (SSL), Wired Equivalent Privacy (WEP), Wi-Fi Protected Access (WPA), and Transport Layer Security (TLS). RC4 uses an XOR function to combine a pseudorandomly generated stream of bits (the keystream) with the plaintext and produce ciphertext. To generate the keystream, the cipher uses a secret internal state composed of two parts: a permutation of all 256 possible bytes and two 8-bit index pointers. The permutation is initialized with a variable-length key, and then a pseudorandom stream of bits is generated with a PRNG. RC4 is not recommended for new applications because of its weak use of keys and lack of nonces, but it is still used because of its speed and simplicity.

In addition to RC4, Rivest created other ciphers in the RC family. Notably, RC6, a block cipher, was developed for the Advanced Encryption Standard competition and selected as a finalist. It was not the winning algorithm, but it is still available for applications that require encoding functions. It is not free, however, because it is a patented algorithm.

Rijndael **Rijndael** (pronounced like *raindoll*) is the encryption algorithm incorporated into the Advanced Encryption Standard (AES). It is a block cipher composed of 10 to 14 rounds of S-box and XOR functions, but it does not use a Feistel network. This symmetric algorithm specifies how to use 128-bit, 192-bit, or 256-bit keys on 128-bit, 192-bit, or 256-bit blocks. It applies 10 rounds for 128-bit keys, 12 rounds for 192-bit keys, and 14 rounds for 256-bit keys. The standard AES implementation of Rijndael has a fixed block size of 128.

The name *Rijndael* comes from the names of its creators, the prominent cryptographers Vincent Rijmen and Joan Daemen. The names *Rijndael* and *AES* are used interchangeably, but they are not the same. Rijndael is the algorithm and AES is the implementation.

Rivest, Shamir, Adelman Rivest, Shamir, and Adelman developed RSA (the first letter of each creator's surname) for public key encryption. RSA uses a public key that is freely shared and a private key that is kept secret. If RSA is used with long enough keys and is

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

kept updated, it is believed to be secure. RSA is widely used in e-commerce protocols and is the default encryption and signing scheme for X.509 certificates.

RSA keys are usually 1024 to 2048 bits, although larger and smaller keys are supported. Keys that are smaller than 512 bits can be broken quickly with modern computing resources. Even the security of 1024-bit keys has been called into question since 2003. Therefore, experts recommend keys that are 2048 bits or longer. As with any asymmetrical encryption system, key exchange is a challenge. Key distribution must be protected against man-in-the-middle attacks, typically by using Public-key Infrastructure (PKI) components, such as digital certificates, which are signed and can be authenticated with hashing functions.

RSA is also vulnerable to timing attacks, adaptive chosen ciphertext attacks, and branch prediction analysis. Timing attacks are a type of side channel attack, which is discussed later in this chapter. Adaptive chosen ciphertext attacks exploit flaws in the public key cryptography standard (PKCS) scheme to recover session keys; using newer PKCS versions with secure padding schemes is recommended. (You learn about PKCS later in this chapter.) Branch prediction analysis takes advantage of processors that use a predictor to determine whether a conditional branch is likely to be taken in a program's control flow. The attack uses a spy process to determine the private key statistically.

Hashing Algorithms

Hashing algorithms are sets of instructions applied to variable-length input (the message) that generate a fixed-length message digest representing the input. The message digest is used for comparison to ensure message integrity. Hashing algorithms do not provide confidentiality because they do not encrypt the message contents in a way that can be decrypted, but they do provide verification that a message has not been altered. Remember that hashing algorithms are mathematical formulas, and hashing functions are the process the computer uses to generate a hash value.

When a message with a message digest is received, the hashing algorithm is run against the contents again. If the values match, the message is considered unaltered. If the values do not match, the message might have been tampered with or corrupted during transmission. Several hashing algorithms are available, but the most common are Message Digest 5 and Secure Hash Algorithm. Both are discussed in the following sections.

Message Digest 5

Ronald Rivest devised **Message Digest 5 (MD5)** in 1991 as a replacement for MD4, which was not secure. It is used in many cryptographic applications, such as digital signatures and virtual private networks (VPNs). MD5 makes only one pass on data and generates a 128-bit hash value that is displayed as a 32-character hexadecimal number.

Methods have emerged that make it easy to generate collisions in MD5. A collision occurs when computing the MD5 algorithm with two different initialization vectors produces the same hash value. For a hashing algorithm to be considered secure, three conditions must be true: No hash should be usable to determine the original input, no hashing algorithm should be run on the same input and produce different hashes, and a hashing algorithm should not be run on two different inputs and produce the same hash (collision). Collisions can be used to determine the plaintext, and because MD5 is commonly used for password storage, the possibility of deciphering plaintext is a serious security risk. Another risk with MD5 involves the use of rainbow tables: precompiled lookup tables of possible hash-plaintext combinations that

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

are posted on the Internet and that make successful brute-force attacks more likely. They can also be used to generate collisions and make password cracking easier. If one plaintext string generates the same hash value as another, the attacker can use the known information to determine the unknown values. MD5's reputation suffered in 2008 when security researchers demonstrated that, by exploiting MD5 collisions, they could create rogue digital certificates—the electronic documents upon which secure e-commerce is based.

Secure Hash Algorithm The National Security Agency designed **Secure Hash Algorithm** (**SHA**) as a successor to MD5, and it is approved for federal government use. SHA version 1 (SHA-1) is used in many cryptographic applications that require checking message integrity, including SSL, SSH, and IPsec.

The SHA standards include five algorithms: SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512. SHA-1 is the most commonly used algorithm, and the other four are sometimes referred to collectively as SHA-2. SHA-1 generates a message digest of 160 bits, and the number added to the names of the other four algorithms denotes the message digest's length. Table 5-1 summarizes SHA message digest lengths and block sizes, and how many rounds of computation are performed.

Algorithm	Message digest length	Block size	Rounds of computation
SHA-1	160	512	80
SHA-224 and SHA-256	224/256	512	64
SHA-384 and SHA-512	384/512	1024	80

© Cengage Learning 2014

Table 5-1 Summary of SHA algorithms

Methods have been devised to decipher a message's original text based on hash values. The U.S. government became so concerned about SHA-1's vulnerabilities that it was phased out in favor of SHA-2 in 2010. As this book was going to press, a competition was being held to choose the new algorithm that would become SHA-3. However, although the effects of SHA-1 attacks in a lab environment are severe, no significant attack methods outside the lab environment have proved successful in compromising SHA-1.

Until SHA-3 is selected, most cryptographic applications continue to use SHA-1 and sometimes MD5, but storing passwords is still a major concern. In some systems, the password is not saved in the file system; instead, a hashed value of the password is stored. When the system carries out password validation for authentication, it takes user input and generates a hash value to compare with the stored value.

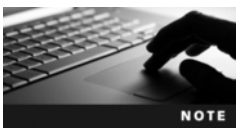

One way to reduce the vulnerabilities of hashing algorithms is to add a salt to the plaintext before hashing. Another method called key strengthening involves applying the hashing function more than once. No method of reinforcement is completely secure, but these methods can help improve the security of MD5 and SHA-1.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Message Authentication Code

Message Authentication Code (**MAC**), also known as Message Integrity Check (MIC), uses a shared secret key that is agreed on by the sender and receiver in the verification process to generate a MAC tag for a message. A MAC tag is like an enhanced message digest. The shared secret key adds a measure of security to the hashing algorithm.

Do not confuse this abbreviation with Media Access Control (MAC), a common abbreviation in networking documentation.

The message and MAC tag are sent to the receiver. The key is also sent to the receiver securely; this key is usually sent separately from the message. The receiver goes through the same process of using the transmitted message and key to generate a MAC tag, and compares this tag with the one received in the message to confirm the message's integrity and authenticity. The verification process is protected by secure communication of the key, which ensures that the sender and receiver generate the same MAC tag from the message. Figure 5-3 illustrates this process.

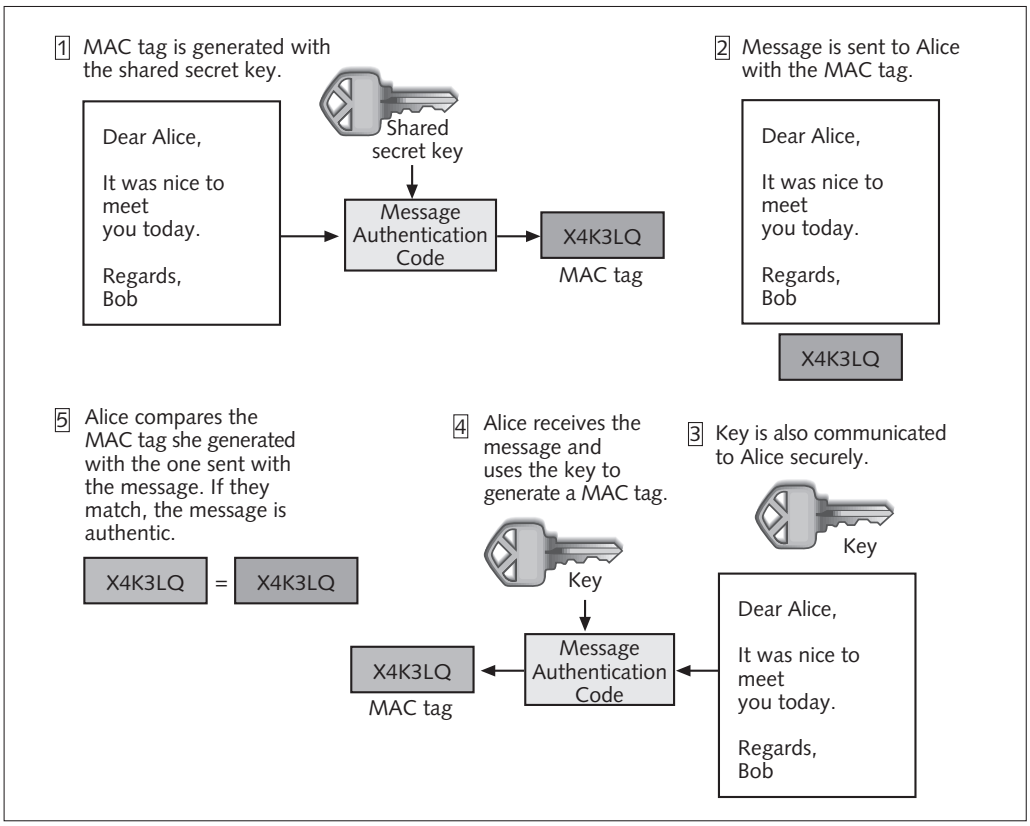

© Cengage Learning 2014

Figure 5-3 The MAC process

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

As with symmetric cryptography, MAC uses a single key to verify message integrity, so the challenge is key management—how to communicate the secret key that the sender and receiver use securely. If communication of the secret key is compromised, an attacker could forge a message between the sender and receiver without MAC detecting it. MAC can still be used, as long as care is taken to secure communication of the secret key with encryption.

Digital Signatures

Digital signatures use hashing algorithms with asymmetric encryption to produce a method for verifying message integrity and nonrepudiation. Nonrepudiation means ensuring that participants in a message exchange cannot deny their roles in the process. This process is explained in the following example and illustrated in Figure 5-4.

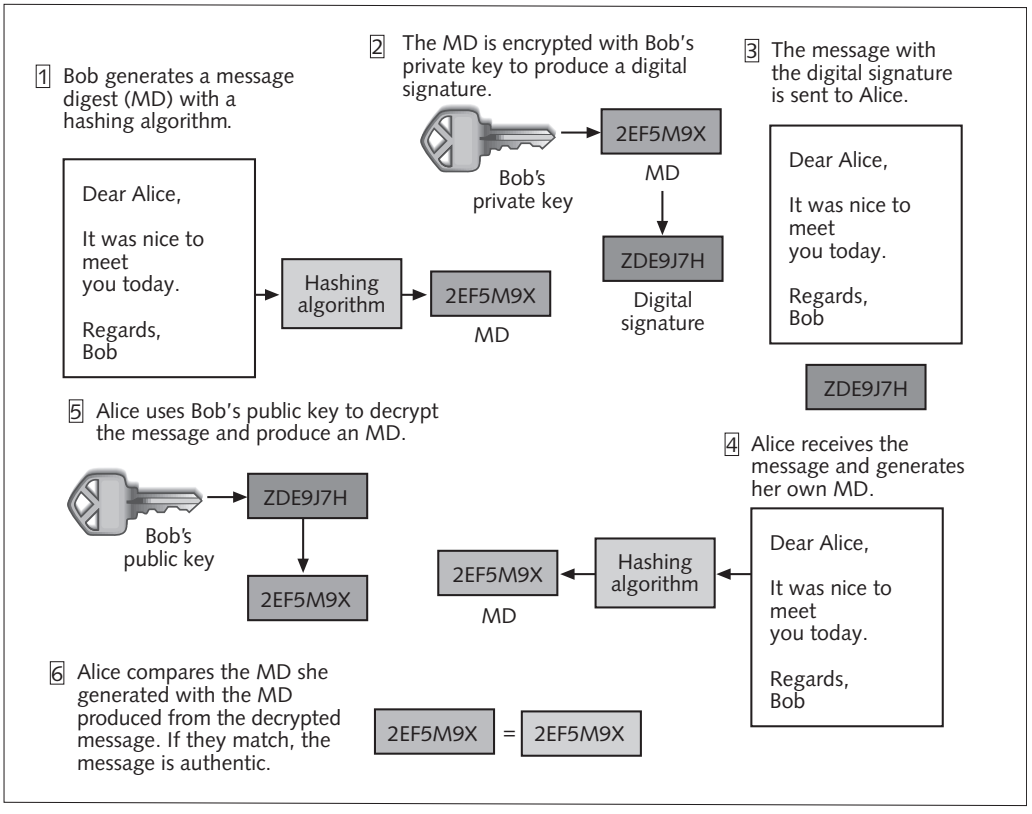

© Cengage Learning 2014

Figure 5-4 The digital signature process

Bob wants to send Alice a digitally signed message. He goes to the organization's directory to find Alice's public key certificate. The process works as follows:

1. Using a hashing algorithm, a message digest of Bob's message to Alice is calculated.
2. The message digest is encrypted by Bob's private key. The resulting ciphertext is the digital signature of the message.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

3. The digital signature with the message is sent to Alice.
4. Alice runs the message she receives through the same hashing algorithm Bob used to get a message digest.
5. Alice then decrypts the message's digital signature with Bob's public key, which produces the message digest calculated in Step 1.
6. Alice compares the message digest she calculated with the message digest Bob calculated to verify the message's integrity. If the message digests are the same, the message is authentic.

Or is it? Alice had to obtain Bob's public key; how can she be sure that the key actually came from Bob? Both Alice and Bob have to trust the source of the public keys.

Digital signature security vulnerabilities are mostly associated with the IT infrastructure required to support interoperability. When a sender digitally signs a message, the receiver must have access to and trust the same certification authority (CA) that issued the sender's credentials. Therefore, Internet users cannot rely on digital signatures because not all users have the configuration to trust issuing CAs. Interoperability and political issues with recognized CAs among organizations and countries are a barrier to universal acceptance of digitally signed documents and messages. Cases are being tried worldwide to determine whether digital signatures can be recognized as a legally valid method for signing documents.

However, if Bob and Alice both trust the CA, nonrepudiation is achieved in this process because the signature is encrypted by Bob's private key. Bob is the only owner of this key and the only person with access to it. The successful decryption by Bob's public key confirms his identity.

Key Management

The major problem with cryptographic algorithms is secure key exchange. To prevent the compromise of encrypted traffic, cryptographic systems change keys frequently and distribute them to all authorized parties. This process of changing and distributing keys is called **key management**, and is difficult to carry out reliably.

Private Key Exchange

Private key exchange uses a symmetric cryptographic algorithm in the encryption process, in which the same key (also called a "shared key" or a "shared secret") is used to encrypt and decrypt a message. Therefore, the message is only as secure as the shared key. The following is an example of a private key exchange, as shown in Figure 5-5.

1. Bob has a cleartext message for Alice. He uses a shared key to encrypt the message into ciphertext to protect it from unintended readers.
2. Bob sends the encrypted message to Alice.
3. For Alice to be able to read the message, she needs to use the shared key to decrypt the ciphertext. This key must be sent to Alice securely and separately to ensure the message's confidentiality. Sending the key with the message or even on the same unsecure medium (in-band) defeats the purpose of encryption.
4. Alice receives the message and uses the shared key she received from Bob to decrypt it.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

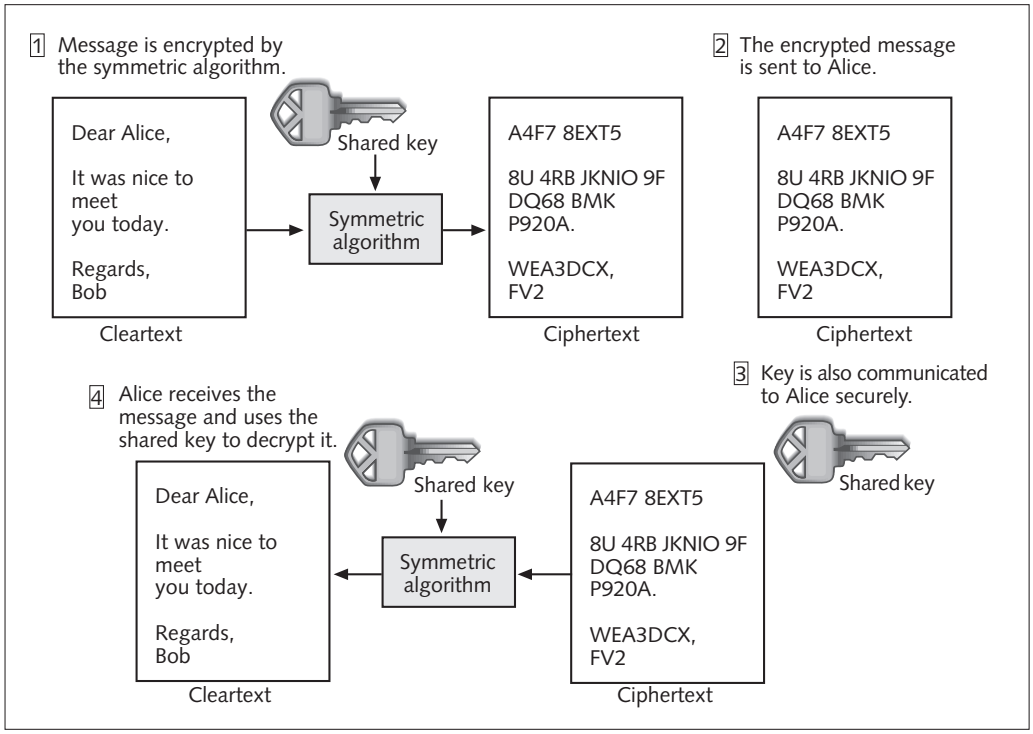

© Cengage Learning 2014

Figure 5-5 The private key exchange process

Public Key Exchange

Public key exchange uses asymmetric cryptography in the encryption process and generates a key pair: Anything encrypted by one key can be decrypted only by the other member of the pair, and vice versa. To achieve key management when a key pair is generated, one is labeled as the public key and the other as the private key. The public key is freely shared—anyone who intercepts it would be able to encrypt messages that the holder of the private key could decrypt. The private key is never shared and is kept secure.

When an encrypted message is sent to the private key owner, the sender encrypts the message with the recipient's public key. The recipient then uses the private key to decrypt the message. Confidentiality is ensured because the private key owner is the only person who can decrypt what the public key has encrypted. Encrypting and decrypting the message with the public key is not possible.

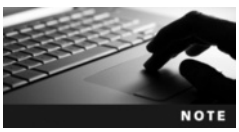

Many computer users use public key exchange in the form of SSL/TLS Web browser encryption and often do not realize it. Typically, asymmetric encryption is not used to provide data confidentiality for the substance of the messages. Because symmetric encryption is so much faster, the two sides use asymmetric encryption to negotiate a shared key securely and then switch to symmetric encryption using the shared key.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The IT infrastructure to support asymmetric cryptography and public key exchange is more complicated than symmetric cryptography and private key exchange. The following list describes components of an asymmetric cryptography system:

- *Certificates*—A certificate is a file that contains information about the user, service, or business entity and the assigned public key. It also contains information about the certification authority and its digital signatures so that users can trust the public key's authenticity.
- *Certification authorities (CAs)*—These organizations issue public and private key pairs to people, services, and organizations. A CA keeps track of issued credentials and manages the revocation of certificates, if needed. A CA also verifies that the public and private keys have been issued legitimately and are trustworthy.
- *Registration authorities (RAs)*—Also called registrars, RAs serve as a front end to users for registering, issuing, and revoking certificates. For security reasons, users rarely contact CAs. Instead, they interact with an RA, which acts as an intermediary in the certificate-issuing process. Before an RA issues public and private keys to a user, identity verification is required to ensure that the certificate reflects the user's identity accurately.
- *Certificate revocation lists (CRLs)*—CAs track and publish listings of invalid certificates. CRLs should be checked to make sure that certificates are legitimate; usually applications perform the check automatically unless users have disabled this feature. If a certificate is listed on a CRL, the cryptography system gives users a warning that the certificate is not valid. Certificates can be considered invalid for many reasons. For example, a user might have been issued a new certificate or taken a new job, so the older certificate needs to be taken out of service. Also, users might no longer be affiliated with the issuing CA because they have been fired.
- *Message digests*—The recipient's message software compares the received message's hash value with the transmitted message's hash value to verify that the message is unchanged.

Putting all these components together, the following process takes place when issuing a certificate for use in public key exchange (see Figure 5-6).

1. Alice applies to an RA for a certificate. The RA works with the CA to issue certificate credentials that contain the public and private keys. In this process, identity verification is carried out via the organization's standard operating procedure for issuing certificates. A copy of Alice's public key certificate is kept in the organization's directory along with issued certificates.
2. Bob wants to send Alice a message. He goes to the organization's directory to get a copy of her public key certificate. Each time the certificate is used, the CA is consulted to make sure the certificate is not listed on the CRL.
3. Bob uses Alice's public key certificate to encrypt the message and then sends the encrypted text to Alice.
4. Alice receives the message and then uses her private key to decrypt it.

Although asymmetric cryptography is commonly used for communicating symmetric keys safely and for security on the Internet, using asymmetric systems for everyday encryption of

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

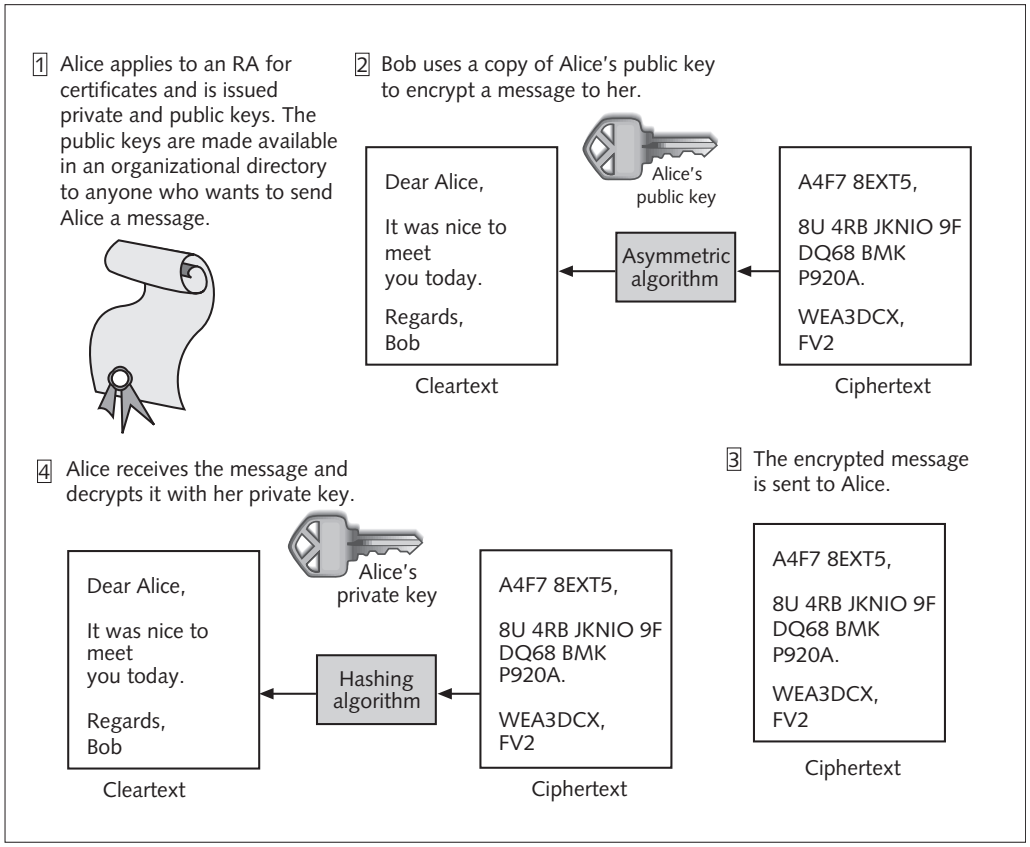

© Cengage Learning 2014

Figure 5-6 The public key exchange process

e-mail and digital signatures has met with mixed success. Asymmetric encryption systems, also known as Public-key Infrastructure (PKI), have been used in the U.S. Department of Defense and large corporations, but the complexity and cost of PKI systems have discouraged many organizations from adopting this technology.

Public Key Cryptography Standards Created by RSA Labs to improve interoperability in public key cryptography, **public key cryptography standards** (**PKCSs**) are not actual industry standards, but they have helped move modern information security cryptography—PKI in particular—toward standardization. Several PKCS designations have even moved into the standards track through the IETF (Internet Engineering Task Force), and others are used as de facto standards. Some PKCS designations have been withdrawn or made obsolete; for example, PKCS #2 and #4 were withdrawn, and PKCS #6 was made obsolete by a newer version of X.509. PKCS #13, Elliptic Curve Cryptography Standard, and #14, Pseudorandom Number Generation, are under development.

X.509 X.509 is an International Telecommunication Union standard for PKI developed in 1998 by the IETF's Public-Key Infrastructure Working Group. X.509 specifies standard

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

formats for public key certificates, a strict hierarchical system for CAs issuing certificates, and standards for CRLs. X.509 certificates use RSA for key generation and encryption and, more recently, SHA-1 hashes to verify the certificate's integrity.

Examining Cryptography Standards

Cryptographic protocols describe how algorithms should be used and specify the security requirements that designs must meet. Often, you hear the algorithm and the protocol that applies it used synonymously—Rijndael and AES, for example. Technically, however, they are not the same. Remember that cryptographic primitives perform a single task reliably but cannot be considered secure when used alone. A **cryptographic protocol** incorporates a detailed description of standardized requirements and guidelines for key generation and management, authentication, encryption, hashing functions, nonrepudiation, and other aspects of ensuring message security.

Reasons for standardizing cryptographic protocols include interoperability, reliability, and scalability, as with most computer standards, but perhaps the most important reason is to attract cryptanalysis attacks. Trying to attract attacks might seem strange, but primitives are proved secure through cryptanalytic efforts to break them. So, the more cryptanalytic efforts are expended on breaking a standard, the better. Standardization in this case provides more rigorous, independent study to verify whether a standard performs as expected and meets its requirements.

The following sections examine major cryptographic standards and some commonly used cryptographic protocols. Keep in mind that many of these systems are also used for remote access and VPN security, wireless security, authentication, and a host of other purposes.

As in most aspects of information security, choosing the most secure or expensive system is not always the best option. Security must be matched to the organization's requirements and plans for the system's use. You should also consider the available budget, resources to handle the system life cycle, and existing infrastructure.

Data Encryption Standard

IBM's **Data Encryption Standard (DES)** algorithm was selected in 1976 as a Federal Information Processing Standard (FIPS). FIPS is an important standard because federal laws mandate its use in certain government projects. DES was an effort to develop a standard way to secure business communication in the United States and was subsequently adopted internationally. It is composed of a 16-round Feistel network with XOR functions, permutation functions, 6×4 S-box functions, and fixed key schedules. Using these primitives, the DES algorithm generates 64 bits of ciphertext from 64 bits of plaintext by using a 56-bit key.

Triple DES

DES is not considered secure for current encryption purposes, mainly because of its 56-bit key size, which is short by today's cryptographic standards. For this reason, a more secure variation called **Triple DES (3DES)** is used now, in which ciphertext goes through three iterations (rounds of encryption). Encrypting with a different key for each iteration increases security. As shown in Figure 5-7, three separate 64-bit keys are used to process the same bit of unencrypted text: The

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

first key encrypts it, the second key decrypts it, and the third key encrypts it again. Using three keys strengthens encryption, but the problem is the time and resources needed to encrypt the information. It takes three times longer than DES, which uses only a single key. Triple DES also requires more processing time, making its use a trade-off.

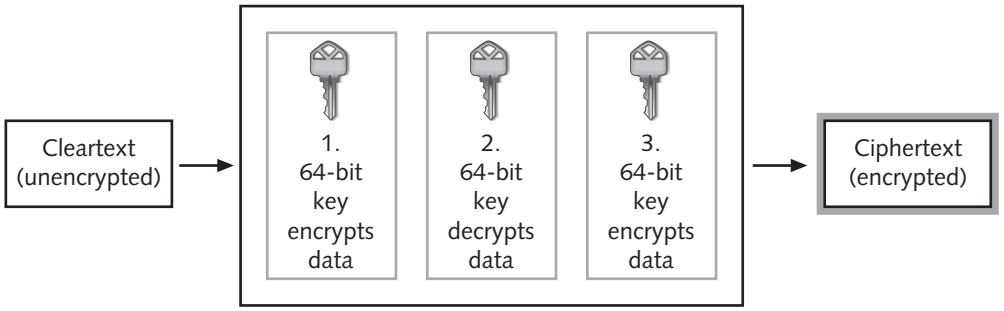

© Cengage Learning 2014

Figure 5-7 3DES encryption

Although most modern computers can use stronger encryption without significant performance degradation, encryption must be selected and managed carefully. Use stronger encryption to ensure confidentiality when needed, and use other forms of security, such as digital certificates, for authentication or integrity. In short, if you do not need to encrypt everything, don't.

Advanced Encryption Standard

The National Institute of Standards and Technology (NIST) approved the **Advanced Encryption Standard** (**AES**) protocol for U.S. government use in 2002. The Rijndael algorithm was selected for AES through a widely publicized international competition. AES is stronger than 3DES and works faster. It is currently the most widely used encryption method, no doubt because of its status as the accepted U.S. government standard and its perceived strength of security. AES is also approved for classified government information.

Although numerous attacks have been attempted against AES, none have been successful as of early 2012. Some of these attempts include side channel attacks, which target the underlying system rather than the cipher implementation. Cryptanalysts have begun to worry that an attacker might be able to stop the computer in the middle of the encryption process, modify the state, and then roll back to the point at which encryption was interrupted. This type of attack may result in the exposure of some keys, but the likelihood that AES-protected data could be read remains extremely remote.

Internet and Web Standards

Because the Internet uses public lines, it is not secure, so methods of keeping messages private while being transmitted over public lines are constantly being developed, tested, broken, and upgraded. The following sections discuss several encryption standards that are used to secure Internet browsing, e-commerce, e-mail, and more.

Secure Shell As you learned in Chapter 4, Secure Shell (SSH) provides authentication and encryption of TCP/IP packets. It works primarily with Linux and UNIX systems,

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

although Windows versions are also available, and it uses public key cryptography. When a client initiates an SSH connection, the two computers exchange keys and negotiate algorithms for authentication and encryption to create a secure connection at the Transport layer. The username and password transmitted to the server are encrypted, as is all data sent subsequently. SSH is available for free with the OpenSSH package (*www.openssh.org*).

Secure Sockets Layer Netscape Communications Corporation developed Secure Sockets Layer (SSL) as a way of enabling Web servers and browsers to exchange encrypted information. SSL is a secure way to transmit data on the Web, including credit card numbers used for online purchases. SSL uses asymmetric keys to start an SSL session and exchange secret keys. After the session is established, SSL uses dynamically generated symmetric keys for the rest of the transfer.

Transport Layer Security **Transport Layer Security** (**TLS**) was designed to provide additional security for Internet communication. Although it is similar to SSL in operation and design, TLS adds some notable improvements. First, it uses a hashed message authentication code (HMAC) that increases security by combining the hashing algorithm with a shared secret key. Both parties must have the same shared secret key to authenticate the data. Second, TLS splits the input data in half, processes each half with a different hashing algorithm, and then recombines them with an XOR function. This method provides protection if one of the algorithms is proved vulnerable.

TLS provides authentication and encryption mechanisms. It supports certificate-based authentication, as SSL does, and the handshake process to establish a session is similar to SSL's. For encryption, TLS can support a variety of symmetric and asymmetric ciphers and several hashing algorithms. TLS uses symmetric keys for bulk encryption and asymmetric keys for authentication and key exchange.

Internet Protocol Security

Internet Protocol Security (**IPsec**) is a set of standard procedures the IETF developed for securing communication on the Internet. For a number of reasons, IPsec has become the standard set of protocols for securing tunneled communication, such as VPNs:

- IPsec works at Layer 3, and therefore provides a type of security that is not available with protocols that work at Layer 2.
- IPsec can encrypt an entire TCP/IP packet.
- IPsec was originally developed for use with IPv6, but it also works with the current IPv4.
- IPsec authenticates the source and destination computer before data is encrypted or transmitted.

Perhaps the biggest advantage of using IPsec is that it is standardized and is supported by a variety of hardware and software devices. In operating systems such as Windows Server 2008 and Windows 7, you can set up an IPsec connection with another IPsec-enabled Windows computer. You add IPsec security policy support as a snap-in to the Microsoft Management Console.

When an IPsec connection is established between two computers, they authenticate one another and then establish the Security Association (SA) settings they use to communicate.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

An SA is a relationship between two or more parties that describes how they use security services to communicate and tracks details of a communication session. SAs are unidirectional, meaning that an SA is set up in each direction of a communication, forming two one-way SAs between parties.

Each IPsec connection can perform encryption, encapsulation, authentication, or a combination of the three. When determining which services to use, the parties in a connection must agree on the details, such as which algorithm to use for encryption. After completing that transaction, the parties must share session keys. These transactions take place in the background. However, in an operating system environment, you need to decide whether IPsec is required for all connections to the host machine or whether the host requests an IPsec connection for computers or other devices that support it. If IPsec is not supported on the client machine, it is not used. With Windows Server 2008 and Windows 7, IPsec is integrated with the Windows Firewall with Advanced Security snap-in.

IPsec Components IPsec's many components provide encryption, encapsulation, key exchange, and authentication. These components include the following:

- *Internet Security Association Key Management Protocol (ISAKMP)*—ISAKMP enables two computers to agree on security settings and establish an SA so that they can exchange keys by using Internet Key Exchange.
- *Internet Key Exchange (IKE)*—This protocol enables computers to exchange keys to make an SA. By default, IKE uses UDP port 500 on both the client and server. Other configurations might use different ports.
- *Oakley*—This protocol enables IPsec to use the Diffie-Hellman encryption algorithm to create keys.
- *IPsecurity Policy Management*—This service runs on Windows computers. It retrieves IPsec security policy settings from Active Directory and applies them to computers in the domain that use IPsec.
- *IPsec driver*—This software handles the tasks of encrypting, authenticating, decrypting, and checking packets.

Suppose that you have configured a VPN connection between two computers and you want the connection to use IPsec. When one IPsec-compliant computer connects to the other, the following process occurs:

1. The IPsec driver and ISAKMP retrieve the IPsec policy settings.
2. ISAKMP negotiates between hosts, based on their policy settings, and establishes an SA between them.
3. The Oakley protocol generates a master key for securing IPsec communication.
4. Based on the security policy established for the session, the IPsec driver monitors, filters, and secures network traffic.

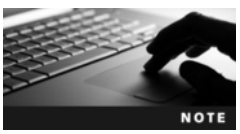

IPsec is not foolproof. For instance, if a computer running IPsec-compliant software has already been compromised, no communication from it can be trusted. IPsec is not a substitute for firewall, antivirus, and IDPS software.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The two core IPsec components are the ones that protect TCP/IP packets exchanged in the connection: Authentication Header and Encapsulating Security Payload. Each is discussed in the following sections.

Authentication Header Authentication Header (AH) is an IPsec component that authenticates TCP/IP packets to ensure data integrity. With AH, packets are signed with a digital signature that tells other IPsec devices that the packet contains accurate IP header information because it originated from an IPsec computer. Digitally signing a packet indicates that it has not been tampered with or that the IP information in the header has not been spoofed. It ensures integrity but not confidentiality.

To authenticate all or part of a packet's contents, AH adds a header calculated by IP packet header and data values, essentially creating a message digest of the packet. Security is achieved by calculating the values with a hashing algorithm and a key known only to parties in the transaction. (Keys are exchanged when the SA is established.) An AH header does not change message contents; it simply adds a field following the IP header. The field contains the computed value of the IP header, except fields that change in transit, such as the time to live (TTL) field. The field also contains data, as shown in Figure 5-8.

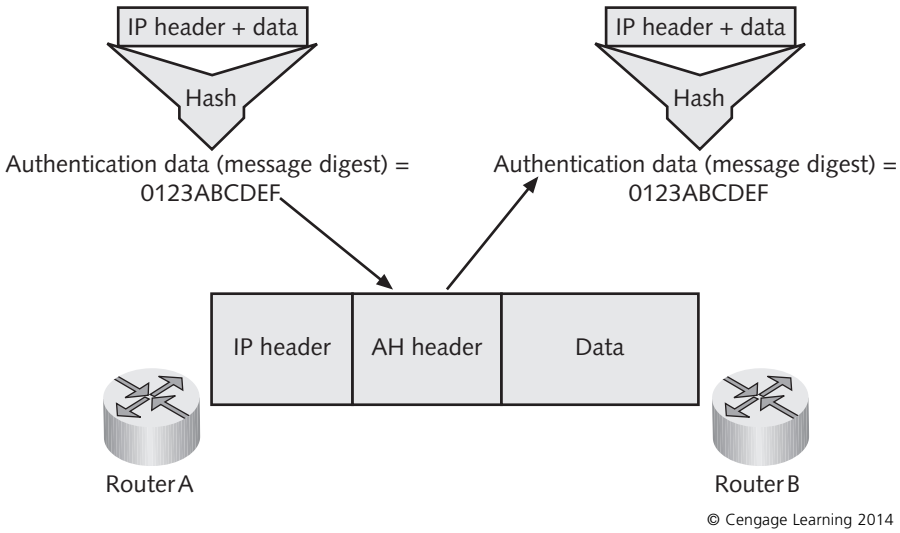

Figure 5-8 AH message exchange

AH works a little differently in the two IPsec modes of tunnel and transport (see Figure 5-9). In tunnel mode, AH authenticates the entire original header and builds a new IP header placed at the front of the packet. The only fields not authenticated by AH in tunnel mode are fields that can change in transit. In transport mode, AH authenticates the data and the original IP header, except fields that change in transit.

AH is an optional protocol. In recent years, it has been implemented less frequently in production environments. While no security problems have been identified in AH, many

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

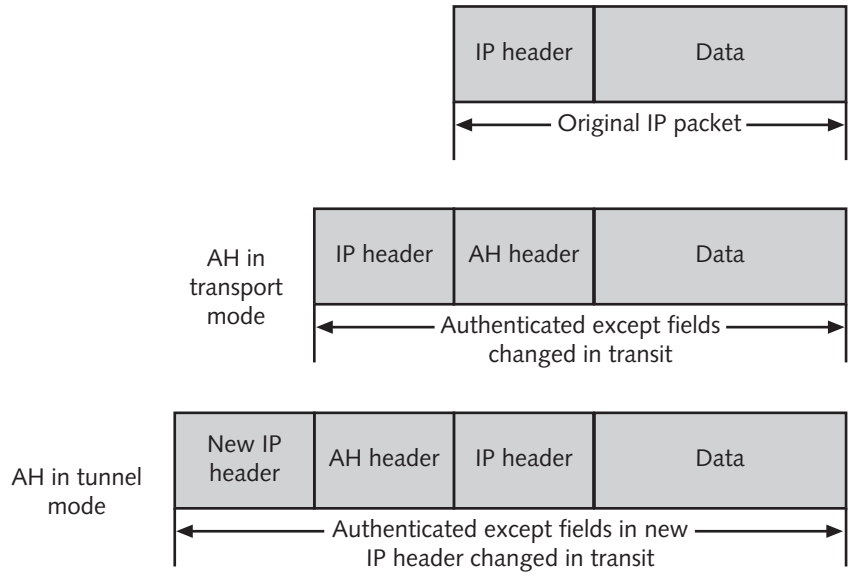

© Cengage Learning 2014

Figure 5-9 AH in tunnel and transport modes

engineers feel that Encapsulating Security Payload (ESP) can perform the same functions in most cases. As of this writing, a draft RFC recommends that AH be retired to decrease complexity of IPsec implementations. Also, ESP using NULL encryption provides the same security in most deployments.

Encapsulating Security Payload AH provides authentication and integrity for messages, but not confidentiality. The confidentiality of data transmitted by using IPsec is ensured by means of ESP. ESP encrypts different parts of a TCP/IP packet, depending on whether IPsec is used in transport or tunnel mode.

In tunnel mode, ESP encrypts both the header and the data portions of a packet. This encryption protects data, but the IP header is encrypted, so the data cannot pass through a firewall that performs network address translation (NAT) because the firewall does not know how to interpret the IP source and destination information in its encrypted form. In transport mode, only the data portion is encrypted. As a result, if the connection passes through a firewall performing NAT, IPsec should be configured to work in transport mode. Figure 5-10 shows the difference the IPsec mode makes to ESP.

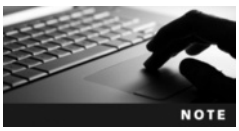

Using AH and ESP together offers additional security, but you might not want to use ESP if another device or application is already providing encryption. ESP, like other forms of encryption, requires substantial processing resources and can slow the rate of data transfer.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

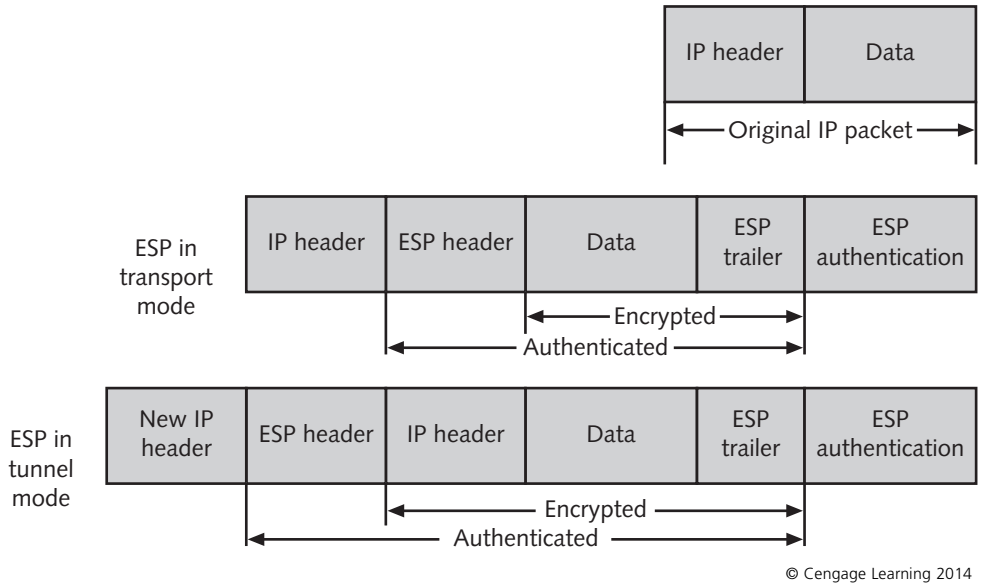

© Cengage Learning 2014

Figure 5-10 ESP in tunnel and transport modes

Modern Cryptanalysis Methods

Many attacks are aimed at cryptographic systems. Some target the encryption algorithm, some exploit flaws in a protocol, and others attack weaknesses in the hardware or software running cryptographic applications. The study of breaking encryption methods is known as **cryptanalysis**. New attacks emerge constantly, so security professionals must keep up to date on threats and countermeasures.

Cryptographic systems present a special challenge, however. Because the expertise to create algorithms or design protocols is based on advanced mathematics, most information security professionals lack the training to address bugs or design flaws. Instead, they must rely on the expertise of mathematicians who design algorithms and agencies, such as the NSA, that test and standardize them. Adhering to guidelines for using and managing cryptographic systems is crucial to maintaining security. The following sections discuss some common attacks against cryptographic systems. You learn the basics of how these attacks work and ways to defend your cryptographic systems from them.

Side Channel Attacks

A side channel attack does not attack the cipher directly; instead, it attacks the underlying systems that leak information, which can be used to compromise the data being stored, processed, or transmitted. These leaks are unintentional signals (emanations) that could expose information being processed. Types of side channel attacks include the following:

- *Timing attacks*—These attacks are based on measuring how much time computations take to perform. For example, a timing attack might watch data move in and out of the CPU or memory. Watching how long it takes to transfer key information can yield clues about key length or eliminating certain key lengths.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *Power monitoring attacks*—These attacks examine hardware's varying power consumption during computations. Watching the power input to the CPU during computations reveals information that can be used to determine the algorithm.
- *Acoustic cryptanalysis*—Similar to power monitoring attacks, acoustic cryptanalysis exploits the sound that computations produce. The current that powers hardware produces heat, which is leaked into the atmosphere. The fluctuations of heating and cooling (thermodynamics) produce low-level acoustic noise that can be examined for clues about the underlying system.
- *Radiation monitoring*—Leaked radiation provides plaintext or other information that can be used to launch an attack. Electrical current fluctuations generate electromagnetic radiation waves, which can occur in patterns. The patterns can be recorded and analyzed to gain information about associated hardware, and sometimes bits of data can be captured.
- *Thermal imaging attack*—If the surface of the CPU can be seen, infrared images can be taken that provide clues about the code.

Side channel attacks rely on emitted information, as in acoustic or radiation monitoring, and relational information, as in timing or power monitoring attacks. Countermeasures against side channel attacks include power conditioning and uninterruptible power supplies to control power fluctuations and emissions, shielding to prevent radiation leakage, and strong physical security to prevent acoustic recorders or other monitoring devices from being installed.

Launching these attacks requires a high level of expertise, which eliminates many potential attackers. As successful attacks emerge, however, automated scripts and instructions invariably appear, making it possible for less knowledgeable attackers to use the techniques.

Passive Attacks

In a passive attack, cryptanalysts simply observe data being transmitted. To gather information, they do not interact with parties exchanging information; they just eavesdrop on transmissions. Detecting this type of attack is difficult because attackers are not transmitting anything. Therefore, countermeasures against passive attacks focus on using strong encryption so that attackers cannot decrypt any data they intercept or crack keys.

Chosen Ciphertext and Chosen Plaintext Attacks

A chosen ciphertext attack selects a captured encrypted message (ciphertext) and decrypts it with an unknown key. This type of attack sometimes uses a decryption oracle, a device that decrypts ciphertext messages the attacker or software has selected. This attack is sometimes called a lunchtime or midnight attack because the attacker gains access to a decryption oracle left unattended during breaks or at night. Chosen ciphertext attacks can be prevented by using the correct cryptographic padding values or redundancy checks.

In a chosen plaintext attack, the attacker can select arbitrary plaintext messages to be encrypted to get the resulting ciphertext messages. Because encryption is carried out in both hardware and software and used in a wide variety of applications, a chosen plaintext attack is often possible.

Public key encryption algorithms that are not randomized are vulnerable to chosen plaintext attacks. Countermeasures are based on randomized encryption, in which a mechanism such as

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

CSPRNGs or randomized padding is used to produce randomized ciphertext messages that cannot be looked up in a rainbow table. Any algorithm that is not vulnerable to chosen plaintext attacks is also considered secure against chosen ciphertext and known plaintext attacks.

XSL Attacks

An XSL attack is a method of block cipher cryptanalysis based on complex mathematical functions (multivariate quadratic equations) that uses an extended sparse linearization algorithm. The researchers who developed the XSL algorithm claim that it can potentially break Rijndael (AES) as well as other block algorithms, such as Camellia and Serpent. This claim is a cause for concern because AES is used in government agencies and many commercial organizations.

Most cryptanalysis methods require an unrealistically high number of known plaintext messages to perform with any effectiveness. XSL attacks require far fewer known plaintext messages to recover a key, which is a major concern. XSL's potential to crack Rijndael is highly debated, but this method warrants further study, and security professionals should monitor new developments.

Random Number Generator Attacks

Modern cryptographic systems require random values for many operations, and hardware or software components that generate or use random numbers can be compromised if attackers can gain access to them via a random number generator attack. They can substitute predictable values and break the entire coding system. These attacks require only a single access to the system, so no further information needs to be sent, as with viruses or worms that obtain a key and then e-mail it back to the attacker. The following major countermeasures can be used against random number generator attacks:

- Combine hardware-generated random numbers with the output of a secure stream cipher. XOR functions are typically used for this method.
- Consider using open-source software for encryption systems. Vendors often do not explain how proprietary products generate random numbers or provide a method to audit the process. Without a way to audit the process, there is no way to assess its security.
- Make sure that physical security for the system is strong.
- Use off-the-shelf hardware for security systems, and do not announce their intended use (such as in online help forums) to prevent potential attackers from knowing what equipment you are using.
- Use a true random source for password generation. Ideally, use a random password/passphrase generator instead of allowing users to choose their own. In practice, this method might be difficult, especially when users must remember their passwords. At a minimum, use the tools provided with your operating system to enforce strong password policies and reduce the possibility of weak passwords.

Related Key Attacks

A related key attack is a form of cryptanalysis in which attackers can observe a cipher's operation by using several different keys. Initial values for these keys are unknown, but a mathematical relationship connecting the keys is known.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Wired Equivalent Privacy (WEP) is an important example of a cryptographic protocol that failed because of related key attacks. Each client in a WEP network uses the same key and the RC4 algorithm. WEP keys must be changed manually, so typically they are not changed often. Attackers can assume that all keys in WEP encryption are related by a known initialization vector (IV). With 24 bits for an IV, only about 17 million keys are possible. This number sounds high, but in practice, WEP's key schedule repeats in a short time. WEP's inherent weaknesses and certain weak keys in RC4 make it easy to recover WEP keys used for encryption. In 2005, the FBI demonstrated that WEP could be broken and the WEP key recovered in less than three minutes.

Preventing related key attacks on WEP networks is not possible, given the protocol's weaknesses. To defend wireless networks against related key attacks and other attacks, use WPA2 or 802.11i for security. To defend wired cryptographic systems against related key attacks or others that exploit weak key schedules, use a cryptographic protocol, such as AES, that incorporates a strong key schedule.

Some older network interface cards (NICs) cannot perform strong encryption methods, so you might need to upgrade NICs to support stronger encryption.

Integral Cryptanalysis

Integral cryptanalysis is applicable to block ciphers that use a substitution-permutation network, including Rijndael, Twofish, and IDEA, among others. This attack uses sets of chosen plaintext messages that share a common constant. Each set of messages shares a constant value, and the remainder of each plaintext message is tried with all possible variables, much like a brute-force attack that checks all possible keys. In integral cryptanalysis, however, only part of the message is tested; the remaining bits are constant. For example, in a set of 256 chosen plaintext messages, each might vary by only 8 bits. In this attack, each set of plaintext messages has an XOR sum of zero, and the corresponding sets of ciphertext messages (generated from the plaintext messages) offer information about the cipher's operation based on the XOR variations.

Differential Cryptanalysis

Differential cryptanalysis applies mainly to block ciphers but can also be used against stream ciphers and hashing functions. Generally, it examines how differences in input affect the output. In block ciphers, it is used to discover where the cipher has nonrandom behavior. Predictable behavior in ciphers results in weaknesses that attackers can use to gain information about the cipher's functions and then recover keys.

Differential cryptanalysis uses pairs of plaintext messages related by a constant difference. By computing differences in the corresponding ciphertext messages (called differentials), attackers might be able to find statistical patterns. Differentials depend on the nature of S-box functions used for encryption, so attackers analyze differentials for each S-box value to look for their frequency of use. This information reveals areas where the cipher displays nonrandom behavior.

Because predictable behavior makes a cipher more vulnerable to being broken, using secure PRNG methods is critical. The goal of cryptographers is to prevent or mask predictable behavior. No algorithm is completely random, but nonrandom functions can be disguised.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Remember the importance of staying informed about the cryptographic system you are using. Sign up for mailing lists and newsletters to keep up to date on emerging threats and new defenses. New versions to correct flaws are released often, so make sure your systems are patched and updated. Another useful place to find information about emerging attacks is hackers' sites. Often, they are the first people to report flaws in cryptographic systems.

Chapter Summary

- Cryptography is the process of converting plaintext into ciphertext by using an encoding function, such as an encryption algorithm or a secret list of substitution characters. Cryptanalysis is the study of breaking encryption methods.
- Cryptographic primitives are modular mathematical functions that are the basic building blocks of cryptographic systems. Each one is designed to perform a specific task reliably. Used alone, however, they cannot ensure adequate security; they must be used with other primitives. Cryptographic primitives include exclusive OR (XOR) functions, permutation functions, substitution box (S-box) functions, Feistel networks, pseudorandom number generators, and hashing functions.
- An encryption algorithm is a set of instructions that provides the encoding function to a cryptographic system or generates output for use in additional operations. Algorithms are classified as block ciphers or stream ciphers. Block ciphers encrypt groups of text, and stream ciphers encrypt cleartext one bit at a time.
- Symmetric algorithms use a shared key in a private key exchange to encrypt a message and decrypt ciphertext. Asymmetric algorithms use two keys for encryption and decryption. Public key exchange is also used to exchange private keys for symmetric algorithms.
- Digital signatures use hashing algorithms with asymmetric encryption for verifying message integrity and for nonrepudiation.
- Public-key Infrastructure (PKI) components include certificates, certification authorities (CAs), registration authorities (RAs), certificate revocation lists (CRLs), and message digests.
- Cryptographic protocols describe how algorithms should be used and specify the security requirements that designs must meet. Common protocols include DES, 3DES, and AES.
- Wireless network cryptographic protocols include WEP, WPA, and 802.11i. For Internet security, SSL, SSH, and TLS are common. IPsec is another major cryptographic protocol used for Internet, VPN, and network security.
- Attacks on cryptographic systems include side channel attacks, passive attacks, chosen ciphertext and chosen plaintext attacks, random number generator attacks, and XSL attacks. In addition, integral and differential cryptanalysis attempt to discover patterns that might reveal nonrandom behaviors in a cipher.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Key Terms

Advanced Encryption Standard (AES) The current U.S. government standard for cryptographic protocols, AES uses the Rijndael algorithm with key sizes of 128, 192, or 256 bits and a fixed block size of 128 bits.

asymmetric algorithm A type of mathematical formula that generates a key pair; ciphertext generated by one key can only be decrypted by the other key.

block cipher A type of encryption algorithm that encrypts groups of cleartext characters.

ciphertext Unreadable text, programs that do not execute, and graphics you cannot view.

cryptanalysis The study of breaking encryption methods. Some common attack methods against cryptographic systems include differential and integral cryptanalysis, random number generator attacks, side channel attacks, and XSL attacks.

cryptographic primitives Modular mathematical functions that perform one task reliably. They form the basic building blocks of modern cryptography.

cryptographic protocol A detailed description that incorporates standardized requirements and guidelines for key generation and management, authentication, encryption, hashing functions, nonrepudiation methods, and other aspects of message security.

cryptography The process of converting plaintext into ciphertext by using an encoding function.

Data Encryption Standard (DES) An older protocol composed of a 16-round Feistel network with XOR functions, permutation functions, 6×4 S-box functions, and fixed key schedules. DES generates 64 bits of ciphertext from 64 bits of plaintext by using a 56-bit key.

digital signature A method of verifying nonrepudiation and integrity in messages.

encryption algorithm A precise set of instructions that provides an encoding function for a cryptographic system or generates output for use in additional operations.

exclusive OR (XOR) function A cryptographic primitive based on binary bit logic and used as a linear mixing function, combining values for use in further computations.

Feistel network A cryptographic primitive that forms the basis of many symmetric algorithms. Feistel networks combine multiple rounds of repeated operations, such as processing cleartext input with XOR functions. A key schedule is used to produce different keys for each round.

hashing algorithms Sets of instructions applied to variable-length input (the message) that generate a fixed-length message digest representing the input. Hashing algorithms do not provide confidentiality because they do not encrypt the message contents, but they do provide verification that a message has not been altered.

hashing functions Processes that a computer runs to verify message integrity by generating a hash value (also known as a message digest), which is a fixed-size string representing the original input's contents. Hashing functions are also used for error detection.

Internet Protocol Security (IPsec) A set of standard procedures that the Internet Engineering Task Force (IETF) developed for enabling secure communication on the Internet.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

key management A way to prevent keys from being discovered and used to decipher encrypted messages. One form of key management is to change keys frequently.

Message Authentication Code (MAC) A hashing algorithm that uses a shared secret key to generate a MAC tag for a message.

Message Digest 5 (MD5) A widely used hashing algorithm that produces a 128-bit hash value displayed as a 32-character hexadecimal number.

permutation functions Bit-shuffling cryptographic primitives that reorder sets of objects randomly.

plaintext Readable text, programs that execute, and graphics you can view.

private key exchange In symmetric cryptography, the same key is used to encrypt and decrypt a message. Public-key Infrastructure is often used for private key exchange.

pseudorandom number generators (PRNGs) Cryptographic primitives used to generate sequences of numbers that approximate random values.

public key cryptography standards (PKCSs) A set of standards that RSA developed to provide standardization guidelines for cryptography. Many of these 15 standards have moved into the IETF standards track.

public key exchange In asymmetric cryptography, two keys are required: the public key and the private key. The public key used to encrypt the message is shared freely. The private key used to decrypt the message is kept secret.

Rijndael The encryption algorithm used in AES; a symmetric block cipher composed of 10 to 14 rounds of S-box and XOR functions. It supports 128-bit, 192-bit, or 256-bit keys and block sizes. Rijndael applies 10 rounds for 128-bit keys, 12 rounds for 192-bit keys, and 14 rounds for 256-bit keys.

Secure Hash Algorithm (SHA) A hashing algorithm that the NSA designed as a replacement for MD5. SHA-1 produces a 160-bit message digest.

stream cipher A type of encryption algorithm that encrypts one bit at a time.

substitution box (S-box) function A cryptographic primitive that transforms a number of input bits into a number of output bits and produces a fixed or dynamic lookup table.

symmetric algorithm A type of mathematical formula in which the key for encrypting cleartext is the same key for decrypting ciphertext.

Transport Layer Security (TLS) A protocol designed to provide additional security for Internet communication. TLS uses a hashed message authentication code (HMAC) to combine the hashing algorithm with a shared secret key. TLS splits input data in half, processes each half with a different hashing algorithm, and recombines them with an XOR function.

Triple DES (3DES) An enhanced variation of DES that uses three 64-bit keys to process data. See also *Data Encryption Standard (DES)*.

X.509 An International Telecommunication Union standard for PKI that specifies standard formats for public key certificates, a strict hierarchical system for CAs issuing certificates, and standards for certificate revocation lists. X.509 certificates use RSA for key generation and encryption, and MD5 hashes to verify the certificate's integrity.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Review Questions

1. Define cryptographic primitives.
2. Which of the following is used as a cryptographic primitive? (Choose all that apply.)
   a. pseudorandom number generators
   b. hashing functions
   c. Feistel networks
   d. side channels
3. Why are cryptographically secure pseudorandom number generators so important to cryptography?
4. What is the block size in the AES implementation of Rijndael?
   a. 128 or 256 bits
   b. 128, 192, or 256 bits
   c. variable
   d. 128 bits
5. Which of the following issues public and private key pairs?
   a. certificate publisher
   b. certification authority
   c. certificate revocation list
   d. certificate store
6. Which of the following is used to check whether a certificate is still valid?
   a. certificate revocation list
   b. certification authority
   c. certificate publisher
   d. registration authority
7. Which of the following is a symmetric algorithm that is not considered safe for encryption use?
   a. AES
   b. Diffie-Hellman
   c. DES
   d. RSA
8. In digital signatures, which of the following values is compared to verify a message's integrity?
   a. public key
   b. message digest
   c. private key
   d. certificate

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

9. When using symmetric and asymmetric algorithms to encrypt the same amount of data, which of the following statements is correct?
   a. The symmetric algorithm encrypts data faster than the asymmetric algorithm.
   b. The asymmetric algorithm encrypts data faster than the symmetric algorithm.
   c. The symmetric and asymmetric algorithms work at the same speed to encrypt data.
   d. The faster an asymmetric algorithm works, the stronger its encryption.
10. Which of the following combines a hashed message authentication code with a shared secret key, processes each half of the input data with different hashing algorithms, and recombines them with an XOR function?
    a. SSL
    b. SSH
    c. TLS
    d. WPA
11. Which of the following is a reason that IPsec has become the standard protocol for tunneled communication? (Choose all that apply.)
    a. IPsec is fast and supported universally.
    b. IPsec supports IPv4 and IPv6.
    c. IPsec is implemented at Layer 2.
    d. IPsec can encrypt the entire packet.
12. Which of the following components enables IPsec to use Diffie-Hellman to create keys?
    a. Internet Key Exchange
    b. Internet Security Association Key Management Protocol
    c. Oakley
    d. IPsec driver
13. Authentication Header verifies the integrity of TCP/IP packets by signing them with a digital signature. True or False?
14. In tunnel mode, Encapsulating Security Payload encrypts which of the following?
    a. packet header
    b. data
    c. both the header and the data
    d. neither the header nor the data
15. Which of the following attacks might have the potential to exploit AES?
    a. PRNG
    b. visual cryptanalysis
    c. chosen ciphertext
    d. XSL

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Projects

Hands-On Project 5-1: Using a Symmetric Algorithm in Word 2010

Time Required: 5 minutes

Objective: Explore the simple encryption used in Word 2010.

Description: This project requires Microsoft Word 2010. In this project, you use Word 2010 to encrypt a document's contents and view the resulting ciphertext. You also determine which cryptographic protocols are available in your operating system.

1. Log on either to the Windows Server 2008 or Windows 7 computer. In a new Word document, type **The quick brown fox jumped over the lazy dog.**
2. Click **File**. The Info button should be highlighted in the left frame. If not, click **Info**.
3. Click the **Protect Document** button in the Permissions section within the middle frame. Click **Encrypt with Password**.
4. In the Encrypt Document window, type **Pa$$word** in the Password text box, and click **OK**. Reenter the password in the next text box, and click **OK**.
5. Click **Save As**, and save the document as **Encryption.docx** to your desktop.
6. Click the **Start** button, click **All Programs**, click **Accessories**, and click **Notepad**.
7. From the Notepad **File** menu, click **Open**. By default, the extensions menu lists Text Documents (*.txt). Click the down arrow, and select **All Files** (*.*) from the menu. Browse to the desktop and open **Encryption.docx**.
8. Now you can see the encrypted version of the file. However, not everything is encrypted. In the metadata of the file, you can see details about the cryptographic protocols used (see Figure 5-11).

```
 ⌟                                                                    ⌟ ⌟ @
  <?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<encryption
xmlns="http://schemas.microsoft.com/office/2006/encryption"
xmlns:p="http://schemas.microsoft.com/office/2006/keyEncryptor/pa
ssword"><keyData saltSize="16" blockSize="16" keyBits="128"
hashSize="20" cipherAlgorithm="AES"
cipherChaining="ChainingModeCBC" hashAlgorithm="SHA1"
saltValue="vCCaOPPQ7N/kaTy3z6cDWA=="/><dataIntegrity
encryptedHmacKey="fr86SVRv2ihwTZzjUYVdc9ERqyyyyyyyyyyyyyyyyyyyyyy
yyyyyyyyyyyyyyyyyyyy
```

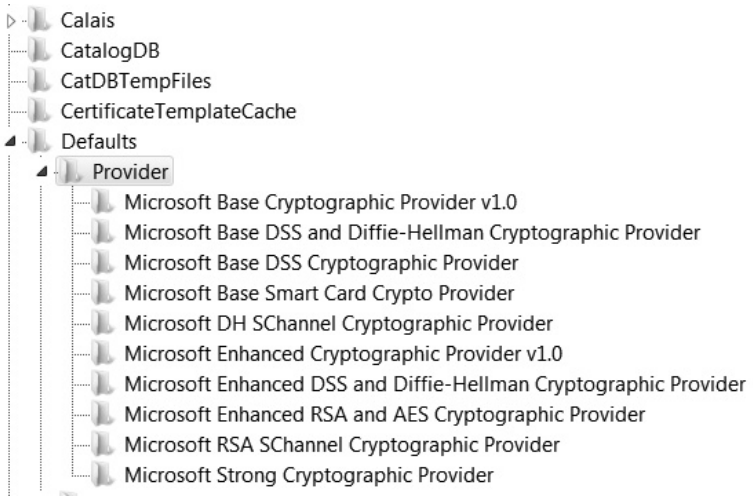

Source: Used with permission of Microsoft Corporation

Figure 5-12 Cryptographic protocols in the Registry

Hands-On Project 5-2: Using TrueCrypt for Virtual Volume Encryption

Time Required: 20 minutes

Objective: Use an open-source encryption tool on a designated volume.

Description: In this project, you download and install TrueCrypt, an open-source encryption tool. A wizard steps you through the process of encrypting a volume on your hard drive.

1. Start your Web browser, go to **www.truecrypt.org**, click **Downloads**, and download the most recent version of TrueCrypt for Windows.
2. Double-click the downloaded file, and run the installation program. If necessary, click **Yes** in the UAC window. In the TrueCrypt License window, click **I accept the license terms**, and then click **Next**. In the Wizard Mode window, accept the default **Install** selection, and click **Next**. In the Setup Options window, accept the default settings, and click **Install**. In the installation completion window, click **OK** and then click **Finish**. After TrueCrypt is installed, click **Yes** if you want to view a short tutorial.
3. Start TrueCrypt, which opens to the main window shown in Figure 5-13. Click the **Create Volume** button to start the volume creation wizard.
4. Verify that **Create an encrypted file container** is selected, and click **Next**.
5. Verify that **Standard TrueCrypt volume** is selected, and click **Next**.
6. Click the **Select File** button to create a file-based volume. Browse to your desktop. In the File name text box, type **Secret** and click **Save**. In the Volume Location window, click **Next**.
7. In the Encryption Options window, click the different encryption algorithms, and read the available information. When you are finished, click **AES**. In the Hash Algorithm section, review the available types of hashing algorithms. Click **SHA-512**, and then click **Next**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

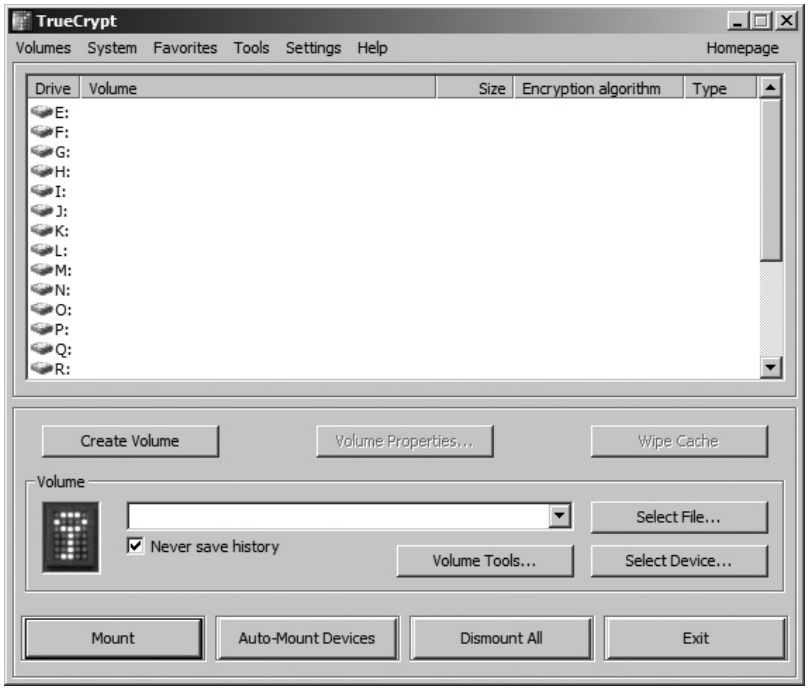

Source: TrueCrypt

Figure 5-13 The main window in TrueCrypt

8. In the next window, enter **10** MB as the container size for your encryption volume, and then click **Next**.
9. In the Volume Password window, type **Pa$$word** in each text box, and click **Next**. Read the warning and then click **Yes**.
10. In the Volume Format Options window, verify that the Filesystem is **FAT** and that the Cluster is set to **Default**. Click **Format** to create the volume. When you see confirmation that the volume has been created successfully, click **OK** and then click **Exit**.
11. The TrueCrypt main window should be open; if not, start TrueCrypt again. In the main window, select an unused volume drive letter to mount the volume. Make sure that the location and filename you selected in Step 6 are displayed next to the blue Volume logo at the bottom. If they are not, click the **Select File** button to find the file you created.
12. Click the **Mount** button. In the next window, enter the password you created in Step 9, and click **OK** to mount the volume.
13. In Windows Explorer, navigate to the drive letter you designated in Step 11, and view the encrypted volume. If you store information in this volume, it is encrypted automatically.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Project 5-3: Using the Cipher Command

Time Required: 25 minutes

Objective: Use the Encrypting File System to encrypt directories and files.

Description: In this project, you use the cipher command to investigate the functions of the Encrypting File System (EFS) in Windows 7 or Windows Server 2008. EFS requires that users have an EFS digital certificate for encrypting and decrypting files and directories. Users who do not have an EFS certificate have one created automatically the first time they encrypt a file. Users can choose from certificates to allow other users to access encrypted files.

1. Log on either to the Windows Server 2008 or Windows 7 computer as the Team*x* administrator.
2. Open a command prompt, type **cd ** and press **Enter**.
3. Create a directory called Private by typing **mkdir Private** and pressing **Enter**. Type **cd Private** and press **Enter** to change to the Private directory.
4. Create a file called SecretFile.txt by typing **copy con SecretFile.txt** and pressing **Enter**. Note that you have not returned to the command prompt; you are still in the file-creating mode.
5. Type **This is a secret file**. Press **Enter**, press **Ctrl+Z**, and then press **Enter** to end the file. Type **dir** and press **Enter** to verify that the file exists. Type **type SecretFile.txt** to verify the contents of the file.
6. Type **cd ..** and press **Enter** to return to the parent directory (C:\). Type **cipher** and press **Enter** to display the encryption status for the contents of C:\. Notice that no files or directories are encrypted, as indicated by the "U" (for "unencrypted") that precedes each listing.
7. Type **cipher /?** and press **Enter** to see the cipher command help file.
8. Type **cipher /E Private** and press **Enter** to encrypt the Private directory. Type **cipher** and press **Enter** to verify that the Private directory has been encrypted.
9. Type **type \Private\SecretFile.txt** to verify that you can still access the file.
10. Create a new file in the Private directory named SecretFile2.txt. Type **copy con C:\Private\SecretFile2.txt** and press **Enter**. Type **This is also a secret file**. Press **Enter**, press **Ctrl+Z**, and then press **Enter** to end the file. Verify that the SecretFile2.txt file exists and that its content is "This is also a secret file."
11. Type **cd Private** and press **Enter**. Type **cipher** and press **Enter** to see the encryption status of the two files.
12. On the Windows Server 2008 computer, create a domain user named **Shifty Smith** with a logon name of **ssmith** and a password of **Pa$$word**. Log off either computer, and then log on as Shifty Smith.
13. Open a command prompt, and change to the C:\Private directory. Enter the **type** command in an attempt to display the contents of SecretFile.txt and SecretFile2.txt. You should be denied access to SecretFile2.txt.
14. Log out and then log on as the Team*x* administrator. Open a command prompt, and change to the C:\Private directory.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

15. Type **cipher /C** and press **Enter** to see details about the encrypted file, including which users can decrypt the file.
16. Click **Start**, click **Computer**, double-click **Local Disk (C:)**, and double-click **Private**. (Note that Private is shown in green to indicate that it is encrypted.) Right-click **SecretFile2.txt**, click **Properties**, click **Advanced**, click **Details**, and click **Add**. Note that Shifty Smith does not have an EFS digital certificate to select, which means that you cannot give Shifty the ability to see the encrypted file.
17. Log off and then log on as Shifty Smith. Right-click the desktop, click **New**, click **Text Document**, and name the document **Shifty**. Right-click the document, click **Properties**, click **Advanced**, check the **Encrypt contents to secure data** box, and click **OK**. Click **OK** again. In the Encryption Warning window, click the **Encrypt the file only** option button, and click **OK**. Log out.
18. Log on as the Team*x* administrator. Repeat Step 16; this time you should see Shifty's EFS certificate. Select it and then click **OK** four times. Log out.
19. Log on as Shifty Smith, and repeat Step 13. This time you should have access to SecretFile2.txt. Type **cipher /C** and press **Enter**. Note that Shifty now is listed as having the ability to decrypt SecretFile2.txt.
20. Close all windows and log out of both systems.

Case Projects

Case Project 5-1: Configuring Encryption in Word 2010

The chief information officer (CIO) in your organization has expressed concerns about the use of encryption by employees. Currently, users encrypt Microsoft Word documents using the built-in encryption function. The CIO has been considering a third-party cryptographic product, but needs more information about how Word encryption works before she can make an informed decision.

You have been asked to write a two-page report that explains the extent to which the IT department can configure the cryptographic features of Word 2010. What is the process involved in configuring encryption?

Case Project 5-2: Using DNSSEC

You are employed by an e-commerce company as an assistant to the chief security officer. Your company has received a number of recent complaints from customers who have ordered products but have not received them. In addition, many of these customers have reported being victims of identity theft. Your company's legal department has recommended that all departments study their practices to ensure that they are taking reasonable measures to prevent identity theft. The chief security officer believes that DNS cache poisoning attacks may be responsible for traffic being redirected from your company's Web site to impostor Web sites where customer information is then collected and misused.

You have been asked to prepare a two-page report that explains how DNS cache poisoning attacks work and how the DNSSEC system can help to prevent these attacks.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 6

Wireless Network Fundamentals

After reading this chapter and completing the exercises, you will be able to:

- Explain wireless communication concepts
- Describe radio frequency characteristics relevant to wireless networking

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Wireless networking is used in government and private industry for an array of services. This chapter introduces you to the fundamentals of wireless communication and wireless networking standards. First, you learn how wireless communication works. You then learn about the components of wireless networks. Wireless standards change rapidly, so you learn the basic standards as well as more recent developments. You also learn how radio communication laws affect wireless communication.

Wireless Communications Primer

Wireless networking is any exchange of data between computers and other devices that does not use cables. Wireless networking is used in many situations, including mobile computing, setting up or adding to a network when running physical cabling is difficult or expensive, and setting up temporary business locations. When a business occupies temporary facilities, setting up a wired network does not make sense; a wireless network is easier, faster, and less expensive. When the business moves, all the equipment can be moved and reused easily in a new location.

Wireless networking follows the same general principles of any computer network—sharing information and resources—but it has some major differences. First and most important, wireless networking uses certain types of electromagnetic radiation, specifically **radio frequency (RF) waves** or **infrared (IR) radiation**, to communicate instead of sending electrical signals over physical cabling. In wireless networks, RF waves are used most commonly; infrared is used mainly for communication with peripheral devices, such as printers, and for short-range communication between laptop IR ports.

Understanding electromagnetic radiation is important for knowing the basics of securing a wireless network.

Electromagnetic Radiation

Cosmic objects emit different types of energy, known as **electromagnetic (EM) radiation**. Electromagnetic energy is made of electrical and magnetic fields, and radiation is energy that travels and spreads out as it moves. So, EM radiation is electromagnetic energy traveling as a self-propagating wave and spreading out at the same time.

A **wave** is simply a means of transporting energy from one place to another without physical movement of material. In wave motion, energy is transported by a disturbance that occurs in a distinct repeating pattern. The maximum departure of the wave from the undisturbed state is called the **amplitude**. Figure 6-1 shows basic wave properties and amplitude.

From highest energy to lowest, the EM radiation types are gamma rays, x-rays, ultraviolet light, infrared light, visible light, microwaves, and radio waves. (More precise classifications exist for these forms of radiation, but they are unnecessary for this discussion.) All these radiation types are collectively called the **EM spectrum**, which is expressed in terms of energy (temperature), wavelength, or frequency. **Frequency** is the number of times an event occurs in a specified period. For example, frequency in EM radiation is measured in hertz (Hz)—the number of waves that occur per second. The greater the number of hertz, the higher the frequency is and the more energetic the radiation. Higher frequencies are measured in KHz, MHz, GHz, and so on.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

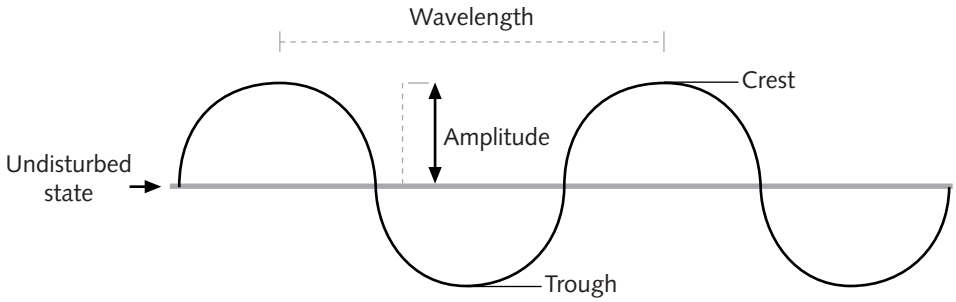

© Cengage Learning 2014

Figure 6-1 Wave properties

Frequency has an inverse relationship with **wavelength**, which is the distance between repeating units (usually the midpoint or crest) of the wave. Frequency is the number of waves per second, and wavelength is the distance between waves. Hotter objects produce higher-frequency, shorter-wavelength radiation than cooler objects.

The atmosphere blocks most of the EM spectrum from striking the earth's surface and allows only visible light, some radio waves, and a small amount of ultraviolet light to pass through. EM radiation at frequencies above 300 GHz, such as microwaves, x-rays, and gamma rays, is lethal.

Infrared Transmissions

Wired networks send information as electrical signals, but **infrared transmissions** use infrared light pulses and require an emitter and a detector. The emitter is usually a laser diode or LED; the detector is sometimes combined with the emitter to detect IR transmissions and generate the corresponding electrical current for wired communication, much like a standard network interface card (NIC) translates data into electrical current (voltage). The intensity of the light pulse in IR indicates the on or off status of each bit of data, just as the presence or absence of voltage indicates the bit status for a wired transmission.

IR communications can be directed or diffused. A **directed IR transmission** requires the emitter and the detector to be pointed directly at one another; a **diffused IR transmission** relies on reflected light that can bounce off walls or other objects. The emitter and detector do not have to be pointed at one another for a diffused IR transmission.

IR wireless has some advantages. It does not interfere with other signals and is not susceptible to interference from them. IR signals cannot pass through walls, but they can pass through glass, so IR transmissions are usually confined to a room. The disadvantages of IR are its limited range, its low speeds of only up to 4 Mbps, and the requirement for direct line of sight or in-the-room conditions. For these reasons, IR is generally used for short-range applications, such as TV remotes, wireless headsets, or short-range data transfers between laptops or handheld computers. Some specialized WLANs (wireless local area networks) use IR, such as in hospitals where RF would interfere with other signals or high-security environments, but in most office settings, IR is not robust enough for WLANs.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Radio Frequency Transmissions

RF wireless transmissions overcome some limitations of infrared, making RF the most commonly used transmission medium for WLANs. RF can pass through walls and other nonmetallic objects and travel great distances, but it is more complex than IR. RF involves transmission ranges, signal modulation, and interference, as you learn in the following sections. Table 6-1 summarizes the most commonly used RF bands. WLANs that use RF operate in the UHF and SHF bands, and primarily use the 2.4 GHz and 5 GHz frequencies.

Band and abbreviation	Frequency	Common uses
Extremely low frequency (ELF; maritime use)	3–30 Hz	Naval and maritime
Super-low frequency (SLF; maritime use)	30–300 Hz	Naval and maritime
Ultra-low frequency (ULF)	300–3000 Hz	Communication in mines
Very low frequency (VLF)	3–30 KHz	Communication with near-surface submarines
Low frequency (LF)	30–300 KHz	Cordless telephones
Medium frequency (MF)	300 KHz–3 MHz	AM radio
High frequency (HF)	3–30 MHz	CB and shortwave radio
Very high frequency (VHF)	30–144 MHz 144–174 MHz 174–328.6 MHz	FM radio TV channels 2–6 Taxi radios TV channels 7–13
Ultra-high frequency (UHF)	328.6–806 MHz 806–960 MHz 960 MHz–2.3 GHz 2.3–2.9 GHz	Public safety Cell phones Air traffic control radar WLANs
Super-high frequency (SHF)	2.9–30 GHz	WLANs
Extremely high frequency (EHF)	30–300 GHz	Radio astronomy
Atmosphere blocks frequencies	Above 300 GHz	N/A

© Cengage Learning 2014

Table 6-1 Common RF bands

Transmission Ranges A wireless network's transmission range varies depending on the standard in use and the environment. Generally, lowering bandwidth (speed of transmission) increases coverage area, while the rate at which a wireless client receives data decreases as the client moves farther away from the transmitter. Higher frequencies usually mean less range.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The range your network needs depends on the site. One pitfall is the urge to increase coverage by having too many access points (APs). An **access point** is an electronic device that connects to a wired network through a cable and can transmit and receive wireless signals through an antenna. Although overall performance and coverage seem to improve with more APs, having stray wireless signals outside your security perimeter is risky. Unlike a wired network with a clearly defined area of operation, wireless signals are transmitted through the air, so enforcing physical security for wireless signals requires careful placement of APs. There is also the risk of **co-channel interference**, in which signals from APs interfere with each other. In an environment with multiple APs, you must be careful when arranging them so that overlapping signals do not share the same channel (frequency).

Walls, metal, and other solid objects impede wireless transmissions, although as a low-frequency radio signal, wireless can pass through walls and other objects to a certain extent. Too many objects or metallic constructions have a detrimental effect on the wireless signal, however. You can overcome this effect by strategically placing APs, adjusting the antennas, or just removing the obstacles. For example, if a wireless workstation is having trouble with its connection, a minor adjustment of an antenna might correct the problem. If this adjustment does not work, you need to examine objects between the wireless NIC (WNIC) antenna and the AP and look for other potential sources of interference.

Interference Radio frequencies can be highly susceptible to interference from electrical storms, solar activity, laser printers, and other forms of EM radiation as well as other RF transmissions. An old microwave oven, for example, can cause major problems for a wireless network because the oven emits RF interference. This interference source is easy to miss if the oven is not in use when you perform an RF sweep to determine sources of interference in your site. (RF sweeps are part of a wireless site survey, as you learn in Chapter 7.) A nearby cordless phone can also cause problems. The factor to remember is transmission power. Devices that use the same or close frequencies compete and interfere with each other, but the device with the higher transmission power continues to operate. The FCC regulates the transmission power of antennas and RF devices.

Wireless communications have a couple of major interference problems. In **multipath**, a signal has more than one path from the transmitter to the receiver. If the signal is reflected or refracted by the atmosphere, ground, or other obstacles, the reflected path can interfere with the direct path. This problem is called **fading**. Co-channel interference, as you learned in the previous section, is another source of interference.

One of the biggest problems in using RF is that concurrent use of the same bands as other entities results in interference that is usually severe enough to disrupt communications. Several RF bands are licensed and can be used only by the licensee on record with the FCC. Unlicensed bands, however, can be used for different purposes by anyone (see Table 6-2).

Radio Frequency Signal Behavior RF signal behavior is characterized by whether a factor contributes to an increase (gain) or decrease (loss) in power. **Gain** is the positive difference in amplitude between two signals and is achieved by magnifying the signal. Although gain and amplitude are sometimes used synonymously, gain is actually a measure of amplification. Gain can occur intentionally when you use a device that provides more power, or

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

unintentionally when the signal bounces off an object and recombines with the original signal, thus increasing its strength.

Unlicensed band	Frequency	Total bandwidth	Common uses
Industrial, Scientific, and Medical (ISM)	902–928 MHz 2.4–2.4835 GHz 5.725–5.85 GHz	234.5 MHz	Cordless phones, WLANs, wireless Public Branch Exchanges (PBXs)
Unlicensed Personal Communications Systems	1910–1930 MHz 2390–2400 MHz	30 MHz	Wireless PBXs
Unlicensed National Information Infrastructure (U-NII)	5.15–5.25 GHz 5.25–5.35 GHz 5.725–5.825 GHz	300 MHz	WLANs, wireless PBXs, campus applications, long outdoor links
Millimeter Wave	59–64 GHz	5 GHz	Home networking applications

© Cengage Learning 2014

Table 6-2 Unlicensed radio frequency bands

Loss, sometimes referred to as attenuation, is the negative difference in amplitude of signals. Sometimes loss is intentional—for example, when you use a device to reduce signal power and comply with regulations on maximum signal strength. Most of the time, however, loss is unintentional. Loss can be the result of many factors, but the following are the most common:

- In **absorption**, certain types of material can absorb RF signals, including wood, concrete, and asphalt.
- In **reflection**, an RF signal can bounce off some materials, which is the opposite of absorption. Objects that cause reflection are usually large in relation to the signal's wavelength and have smooth surfaces, such as walls or buildings. The earth's surface can also reflect signals, and metals are generally reflective as well.
- In **scattering**, small objects and rough textures can disperse the signal. Trees, foliage, rocks, and sand are common objects that cause scattering. Certain weather conditions, such as blowing dust or rain, can also cause scattering.
- In **refraction**, differences in density between air masses over distances can cause problems. When moving through these differences in density, the signal actually bends instead of traveling in a straight line.
- **Diffraction** is similar to refraction, except that the signal bends around an object in its path instead of bending as it moves through the object.
- **Voltage standing wave ratio** (**VSWR**) is caused by differences in equipment rather than external influences, as in the previous examples. If one part of the equipment has a different impedance than others, the signal can be reflected back within the device, and the reflected power can cause the device to burn out (see Figure 6-2).

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

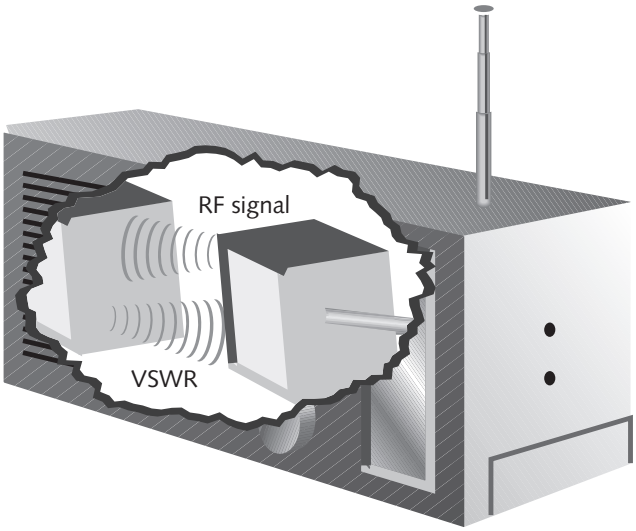

© Cengage Learning 2014

Figure 6-2 Voltage standing wave ratio

Measuring RF Signals Knowing some relatively simple mathematics enables you to understand how to measure RF power and understand the 10s and 3s rules of RF math.

RF power is measured in two ways. The first is on a linear scale using milliwatts (mW). A **watt** is a measure of power or the rate at which work is done. One mW is equal to one-thousandth of one watt of power. In linear-scale systems, everything is relative to a starting point of zero. When calculating gain or loss, however, this linear scale is not helpful because the important consideration is the relation of power measurements to the whole. For example, a loss of 50 mW could be major or minor in relation to the transmission's total power. If the total power is 100 mW, a loss of 50 mW is major. If the total power is 30,000,000 mW, the same loss is minor.

When the transmission power is the point of reference, you get a better measurement of the impact of loss or gain. A 50 mW loss when the power is 100 mW means a 50 percent loss. If the total power is 30,000,000 mW, a loss of 50 mW is only a 0.0001 percent loss. This does not mean that a relative scale is more precise than a linear scale, but that a relative scale provides a better picture of the loss or gain in relation to the whole.

The industry-standard measurements of power are mW and **dBm** (**decibel-milliwatt**). dBm is the reference point that relates the decibel scale to the linear milliwatt scale and specifies that 1 mW = 0 dBm. In other words, instead of comparing the power of a signal to another arbitrary signal, you measure the power of a signal in relation to the reference signal of 1 mW. Thus, a 1 mW signal would be zero dB greater than the reference signal. RF power gains and losses on a relative scale are measured in decibels (dB) instead of mW. **Decibels** measure the difference between two signals. Decibels are a logarithmic method of expressing power and voltage ratios, and are used because they are easier to add and subtract than ratios or percentages when calculating power levels. Decibels are not exact, however; small differences are generally rounded off to make calculations easier.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

By rounding off, you can derive a convenient set of consistent rules: the 10s and 3s rules of RF math, as shown in Table 6-3.

–3 dB equals	half the power in mW
+3 dB equals	double the power in mW
–10 dB equals	one-tenth the power in mW
+10 dB equals	10 times the power in mW

© Cengage Learning 2014

Table 6-3 The 10s and 3s rules of RF math

Thus, you can measure the power of RF signals in dBm and compare power easily. For example, if a signal measures 100 mW, you can convert it to dBm as follows:

> 1 mW times 10 (which is an increase of 10 dB) times 10 (which equals another 10 dB increase) equals 100 mW. Therefore, 100 mW = 20 dBm.

Another signal of 30 mW can be converted to dBm similarly:

> 1 mW times 10 (which is an increase of 10 dB) times 3 (which equals an increase of 3 dB) equals 30 mW. Therefore, 30 mW = 13 dBm.

The other important relative-scale power measurement is **Equivalent Isotropically Radiated Power (EIRP)**. EIRP is the power radiated by a wireless system's antenna, including the intended transmission power and the unintended gain. Although RF gain and loss are measured in decibels, EIRP uses a slightly different measurement known as isotropic decibels (dBi) that applies only to an antenna's gain. The reference point of EIRP is a theoretical antenna operating at 100 percent efficiency.

There are other measurements of RF power, but most measurements for WLANs are in dBm or mW. WLAN measurements are defined in the United States by FCC Part 15.247 and limit the distance that a WLAN can transmit. EIRP measures the gain of the antenna; Transmitter Power Output (TPO) measures the power being delivered to the transmitting antenna and is regulated by the FCC as well.

RF Signaling

RF transmits a **carrier signal**, which can change based on the signal's voltage and direction. A change in voltage does not change the transmission frequency; rather, it determines how long the transmission takes to reach the maximum voltage, fall to the minimum, and return to a neutral charge. In an unmodified RF transmission, however, the signal contains nothing to distinguish between the carrier wave and any other aspects of the signal, such as data. For this reason, changes to the carrier wave are needed to transmit data, video, audio, and other information on the RF signal.

RF data is transmitted as analog or digital signals. An **analog RF signal** is a continuous wave that oscillates between positive and negative voltage, as with an audio or video broadcast. An analog signal has an infinite number of points within the carrier's range. **Digital RF signals** also oscillate between positive and negative voltage, but they are divided into discrete

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

segments or defined states within the carrier's range. When a binary signaling method is used, the range has only two significant states, which are represented by 1s or 0s.

A carrier signal alone cannot carry information, so how can data be transmitted? The answer is **modulation**, which means changing characteristics of the signal. Computers transmit, interpret, and store information in digital format, so an analog RF signal must be converted into digital format. This conversion requires a modulator. Several modulation methods for analog and digital signals are used in wireless networking. You can modify three characteristics of the carrier wave signal to enable it to carry information:

- Height (amplitude)
- Frequency
- Relative starting point of the signal

Analog Modulation The following analog modulation methods (shown in Figure 6-3) are available:

- In **amplitude modulation** (**AM**), the height of the carrier wave is changed so that a higher wave represents a 1 bit and a lower wave represents a 0 bit. (Remember that 1 or 0 is the binary representation of on or off.)
- In **frequency modulation** (**FM**), the number of waves representing one cycle is changed so that the number representing a 1 bit is greater than the number representing a 0 bit.
- In **phase modulation** (**PM**), the cycle's starting point is changed when the bit being transmitted changes from 1 to 0. Instead of the wave oscillating regularly from maximum to minimum when the bit changes, the wave oscillates from minimum to neutral and then back to minimum before returning to neutral and back to maximum.

Digital Modulation Digital modulation techniques still use changes in signal frequency, amplitude, and starting point (phase), but they are superior to analog methods for four reasons:

- More efficient use of bandwidth
- Fewer interference problems
- Error correction that is more compatible with other digital systems
- Less power required to transmit

A digital signal is the same as a binary signal. Three binary signaling techniques can be used with digital modulation:

- *Return-to-zero (RTZ)*—Voltage increases to represent a 1 bit, with zero voltage representing a 0 bit. The voltage for a 1 bit drops back to zero before the end of the bit period. A bit period is the amount of time required to transmit a logical 1 or a logical 0.
- *Non-return-to-zero (NRZ)*—Voltage increases to represent a 1 bit, with zero voltage representing a 0 bit. The voltage for a 1 bit does not drop back to zero before the end of the bit period; it remains high for the duration of the bit period.
- *Polar non-return-to-zero (polar NRZ)*—Voltage increases to represent a 1 bit and drops to negative voltage to represent a 0 bit.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

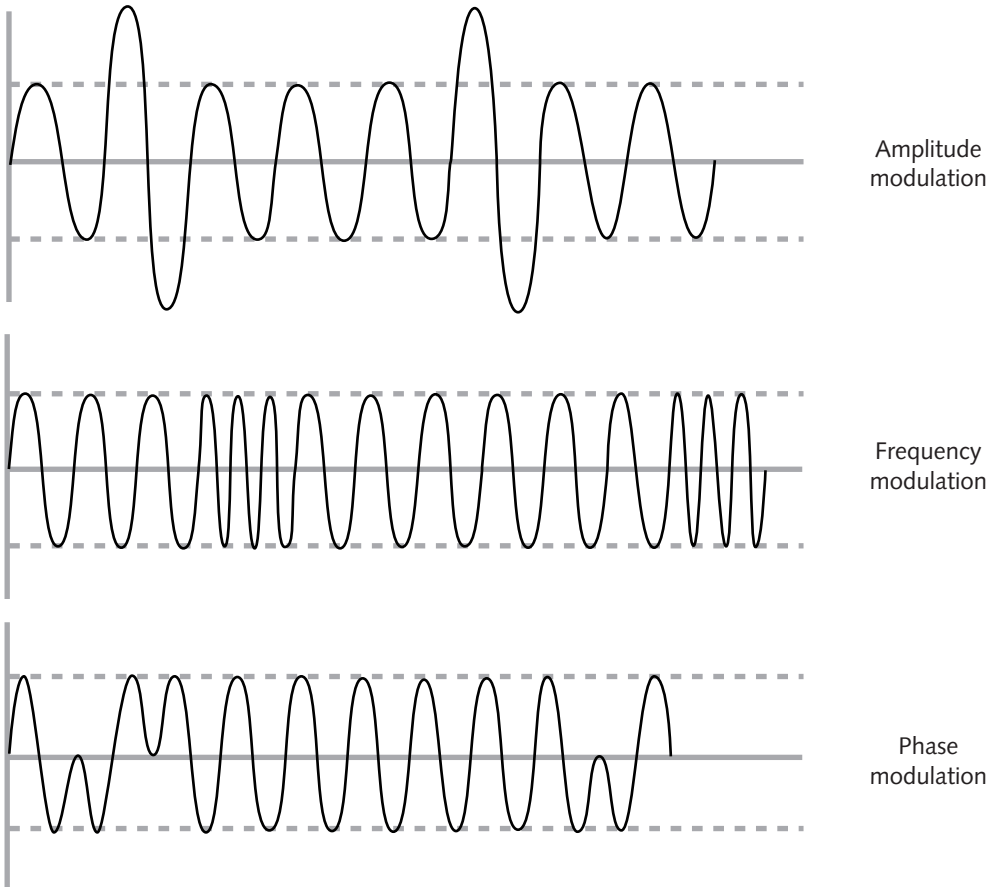

© Cengage Learning 2014

Figure 6-3 Analog modulation techniques

RF signals are by nature **narrowband** transmissions, meaning they transmit on one frequency or a very small frequency range. Narrowband transmissions are particularly vulnerable to interference from any signals operating in or near the frequency being used. Radio and TV stations operate with little disruption on a narrowband frequency because no one else is allowed to use their FCC-licensed frequencies. The following digital signal modulation methods are the most common:

- **Amplitude shift keying** (**ASK**) is a binary modulation technique in which the height of the carrier can be changed to represent a 1 or 0 bit. In ASK, a 1 bit has a carrier signal, and a 0 bit has no signal (in other words, voltage versus no voltage).
- **Frequency shift keying** (**FSK**) is a binary modulation technique in which the carrier signal's frequency is changed to represent a 1 or 0 bit.
- **Phase shift keying** (**PSK**) is a binary modulation technique similar to phase modulation, in which the signal's relative starting point is changed to represent a 1 or 0 bit

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

and changes whenever the bit changes from one to the other. Because PSK is a binary method, the signal starts and stops.

- **Frequency division multiplexing (FDM)** is an RF signaling method in which multiple base signals are modulated on different carrier waves and combined to form a composite signal.

These methods work much like analog modulation techniques, except that digital signals start and stop, and ASK uses NRZ coding instead of both the 1 and 0 bits having carrier signals. A narrowband signal is shown in Figure 6-4.

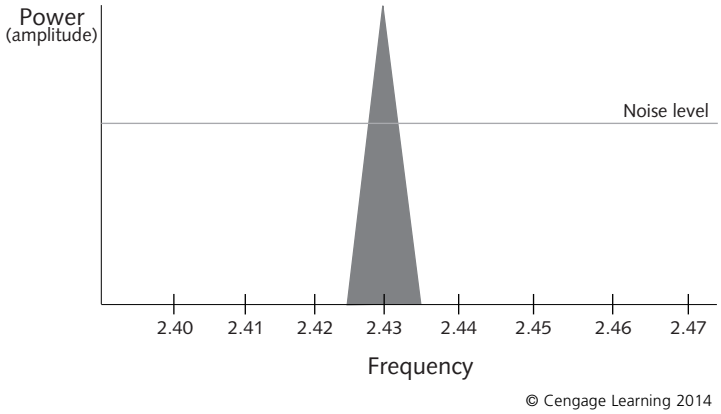

Figure 6-4 Narrowband transmission

Spread Spectrum An alternative to narrowband transmissions is **spread spectrum**, which spreads a signal over a broader portion of the radio band. Aside from reducing interference vulnerability, spread spectrum has some major advantages over narrowband. First, the transmitted signal's bandwidth is much higher than that of the original message. Second, bandwidth is determined by the spreading function and is known only to the transmitter and receiver, which is especially important for security reasons. A spread-spectrum transmission is shown in Figure 6-5.

The more technical definition of spread spectrum is an RF communication system in which the base signal is spread over a larger bandwidth by injecting a higher-frequency signal. The spreading function attaches a key, also called a spreading code or sequence, to the communication channel. The method for inserting the key defines the spread-spectrum technique. The following major methods of spread spectrum are used in wireless networks:

- In **direct sequence spread spectrum (DSSS)**, as shown in Figure 6-6, the key is applied at the data level. The spreading code (key) is the chip sequence used to represent message bits.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

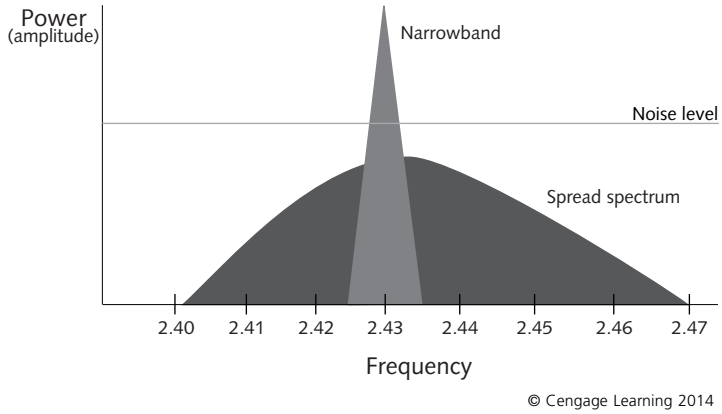

Figure 6-5 Spread-spectrum transmission

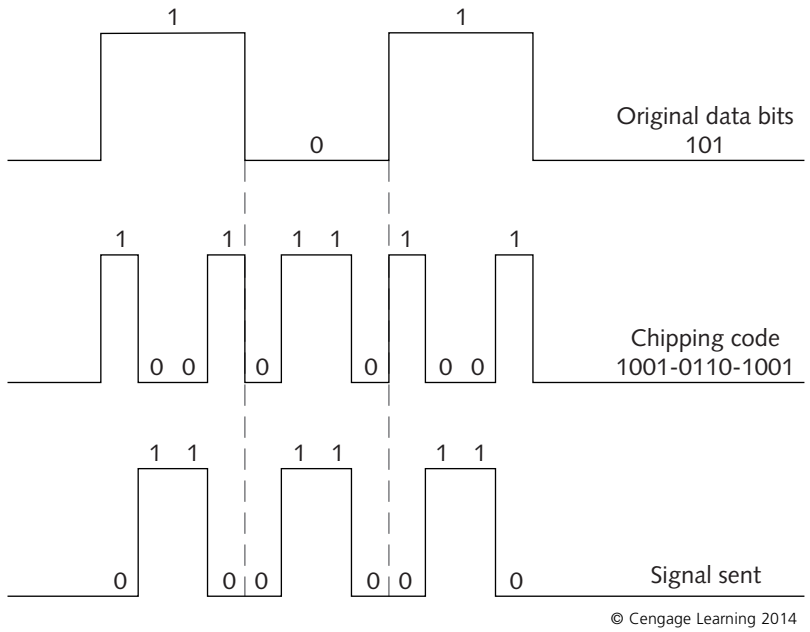

Figure 6-6 DSSS transmission

- In **frequency hopping spread spectrum (FHSS)**, as shown in Figure 6-7, the key is applied at the carrier frequency level. The spreading code (key) is the list of frequencies used for the carrier signal.
- In **orthogonal frequency division multiplexing (OFDM)**, which is technically a modulation technique, a high-speed signal is divided into smaller pieces that are sent simultaneously across many lower-speed channels.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

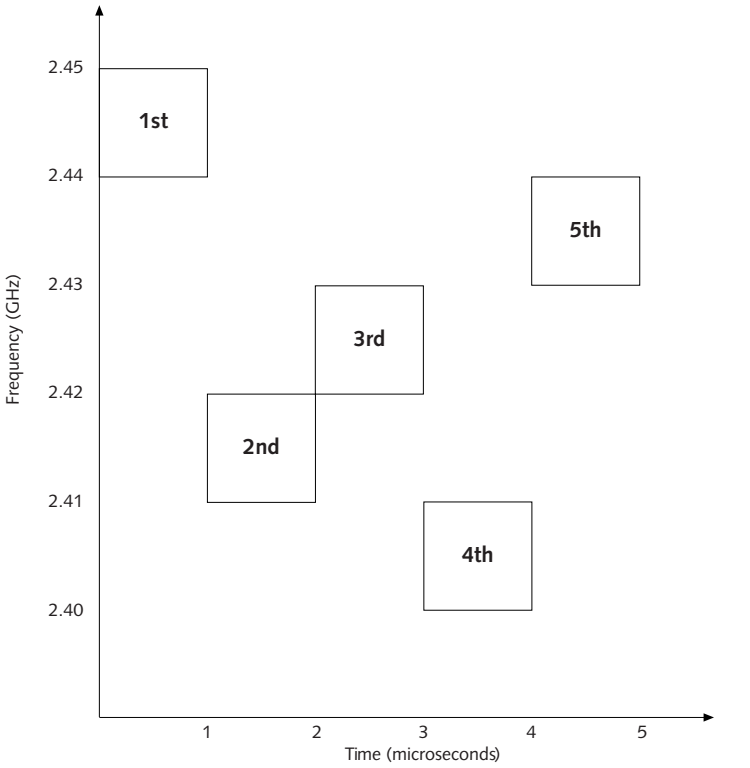

Figure 6-7 FHSS transmission

In DSSS, an expanded redundant **chipping code** is used to transmit each bit. *Chipping code* is the term for the bit pattern. (A single radio bit is often called a chip.) Because the chipping code is applied directly to data entering the carrier modulator, the modulator sees a higher bit rate corresponding to the code sequence's chip rate. Every chip carries data bit information. If interference noise does not affect every chip, information can be recovered, making DSSS less vulnerable to data loss from interference. The downside of DSSS is that spreading sequences tend to be very long, so high bandwidth is needed. However, the signal is not likely to be intercepted because it can barely be detected. To potential eavesdroppers, the signal appears as a slight increase in noise.

In FHSS, the carrier hops frequencies over a wide band according to a sequence defined by the key. The speed of the hops depends on the data rate of the original information. FHSS transmits a short burst on one frequency, then hops, transmits another short burst, and so on until the transmission is completed. The key in FHSS is called the **hopping code**, and it determines the sequence and speed of frequency hops. The major advantages of FHSS are immunity to jamming and interference and low probability of interception, making it fairly secure.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Redundancy of the transmitted message is achieved by the potential to retransmit on different carrier frequencies (in FHSS) and through the message bit on each chip of the spreading code (in DSSS).

With DSSS or FHSS, multiple wireless or RF systems can coexist if different keys (spreading sequences) are used. With the growing demand for higher wireless data rates to support bandwidth-heavy applications, the industry is developing ways to maximize transmission frequency (spread spectrum) and allocate the transmission to multiple users (multiplexing).

Wireless LANs and Their Components

When preparing to secure a WLAN, you need to be familiar with wireless components, topologies, transmission and frequency ranges, and methods of identifying and eliminating interference sources. You learn more about WLANs in Chapter 7.

Wireless communication devices include cell phones, wireless access points, TV remote controls, radios, satellites, medical diagnostic equipment, and more. In WLANs, however, you spend most of your time dealing with networking components that enable wireless communication or that enable the integration of wireless with wired networks.

Wireless NICs

Wireless NICs (WNICs) are the unwired counterpart of standard, wired NICs in networked computers and devices. When a WNIC prepares to transmit, it does the following:

1. Changes the computer's internal data from parallel to serial transmission
2. Divides the data into packets and attaches the address information of the sending and receiving computer
3. Determines where to send the packet (access point, wireless router, and so forth)
4. Transmits the packet

A variety of WNICs are available for desktop computers, laptops, handheld devices, and more. As with standard NICs, WNICs are often added to devices as an on-board feature. Figure 6-8 shows some common add-on types, including (a) a PCI internal NIC, (b) a standalone USB device, and (c) a USB key fob.

Access Points

An access point (AP) consists of the antenna and radio transceiver used to transmit and receive from wireless nodes, an RJ-45 jack for a connection to a wired network, and bridging software to translate communications between wired and wireless devices (see Figure 6-9). APs perform two major functions: acting as the base station for the wireless network segment and serving as the bridge between wired and wireless segments.

The range of an AP varies depending on the type of AP used, the wireless standard, and environmental factors, such as walls, doors, and interference sources. The number of wireless clients an AP can support varies depending on the clients' location and function. Because the

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

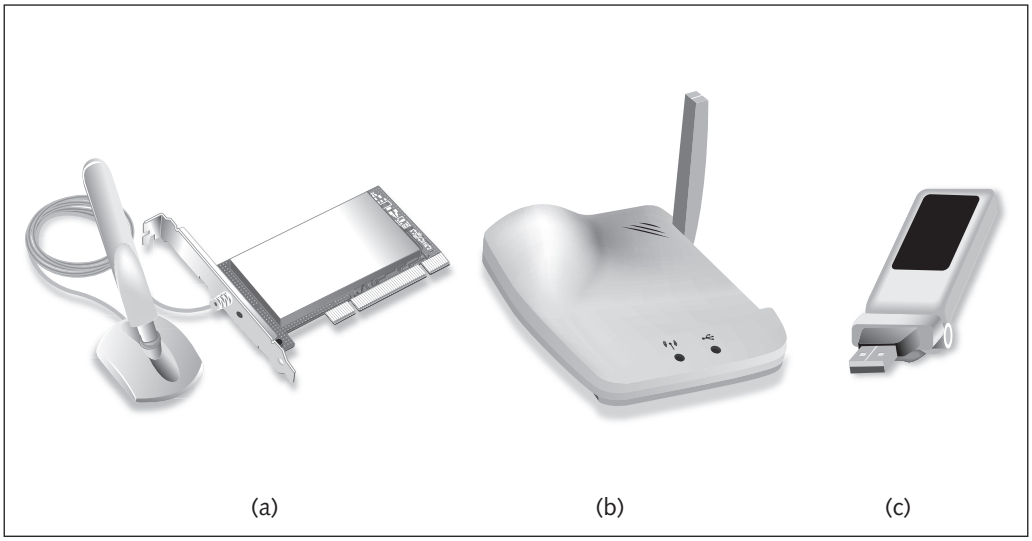

© Cengage Learning 2014

Figure 6-8 Desktop computer WNICs

© Cengage Learning 2014

Figure 6-9 Wireless access point

radio signal is being shared, clients that perform high-bandwidth activities reduce the total number of clients the AP can support.

Generally, APs mounted on desks are subject to interference from solid objects. The preferred placement is on the ceiling or high on a wall, but connecting the AP to an electrical power source can become a problem because outlets are usually close to the floor. The solution is

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Power over Ethernet (**PoE**), in which power for the AP unit is supplied by the unused wire pairs in standard twisted-pair Ethernet cabling, eliminating the need for electrical wiring.

Antennas

RF waves are transmitted and received by an **antenna**. One end of the antenna points into the air and the other end is fixed to the ground directly or via a grounded device. Antennas are generally made of copper wire or another conductive medium; radio waves that strike the wire during transmission or reception cause electrical pressure (voltage) on the wire. The voltage forces the wire to oscillate at the same frequency as the radio waves striking the wire or being emitted. Therefore, RF waves are broadcast by applying an electrical current to the antenna, causing the antenna's wire to oscillate at the same frequency as the radio signal (see Figure 6-10).

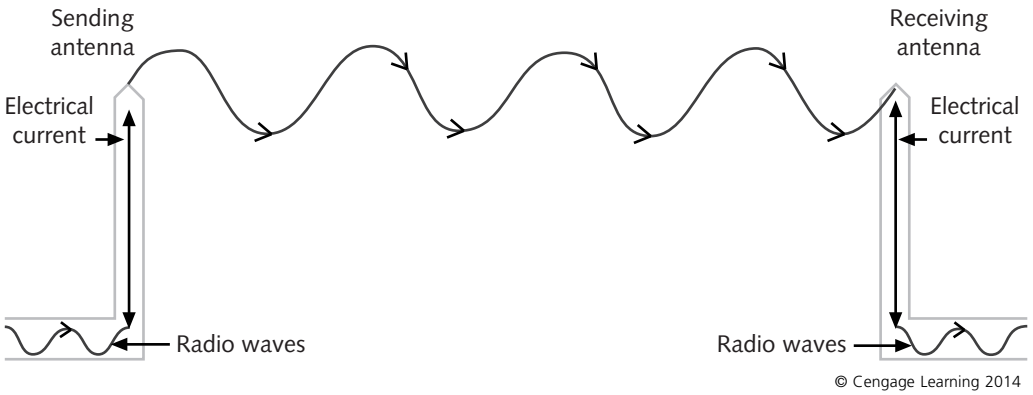

Figure 6-10 Antenna sending and receiving radio signals

EIRP is the measurement of the total power radiated by a wireless system's antenna, but the FCC uses the term **intentional radiator** to describe a device that is designed to generate radio signals, not including those from the antenna. The FCC regulates the power an intentional radiator can generate.

Antennas have the following fundamental characteristics:

- As the frequency gets higher, the wavelength gets smaller, requiring a smaller antenna. Generally, antenna length should be about one-fourth of the wavelength.
- As antenna gain increases, the coverage area narrows.

Other characteristics of RF antenna transmissions include polarization, wave propagation, multipath distortion, the Fresnel zone, and free space path loss. **Polarization** refers to the plane in which radio waves propagate or the orientation of radio waves as they leave the antenna; this orientation is usually vertical or horizontal, but it can also be circular or in any direction. Remember that antennas are wires with electrical currents applied. The current forces the wire to oscillate, generating waves. In basic wave motion, waves follow the plane of their electrical fields, which are parallel to the radiating element. An antenna that is

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

perpendicular to the ground is said to be vertical, and an antenna that is parallel to the ground is said to be horizontal. Therefore, a vertical antenna radiates a signal horizontally across the ground. Antennas that are polarized differently from each other cannot communicate effectively because most of the signal does not strike the receiving antenna. For example, a horizontal antenna and a vertical antenna cannot communicate because the signals radiate toward the tips of the antennas rather than along their length.

Wave propagation is the dispersal pattern of waves as they travel from sending to receiving antennas. Waves can bounce off the earth's ionosphere (sky wave propagation) or follow a direct path to the receiving antenna (RF line of sight). Multipath distortion is not usually a problem with sky wave propagation, but in line-of-sight (LOS) transmission, objects can interfere with the signal, causing refraction or diffraction. The refracted or diffracted signal still reaches the receiver, but the signal might arrive at a different time than the signal propagated directly between antennas. Multipath distorted signals actually reduce the power of a straight-line transmission. One solution to multipath distortion is antenna diversity, which uses multiple antennas, inputs, and receivers.

The **Fresnel zone** (pronounced *fra-nel*) is a series of ellipsoidal shapes in the wave calculated to determine the signal strength. It also identifies potential obstacles and multipath distortion between sending and receiving antennas (see Figure 6-11). As a rule, zone 1 of the Fresnel zone should be 60 percent free of obstacles.

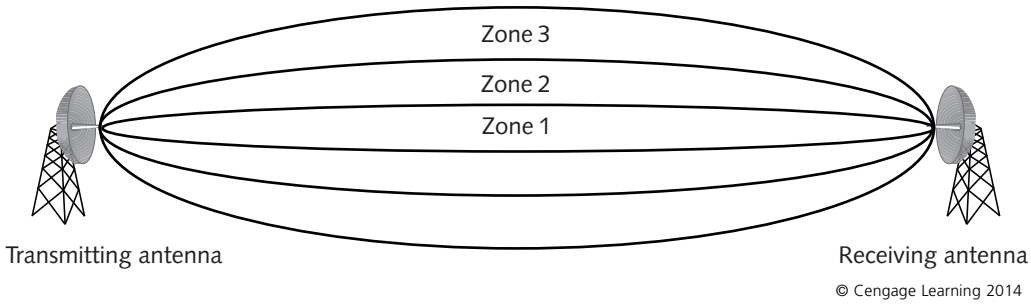

Figure 6-11 The Fresnel zone

As an RF signal travels from the sending antenna, it disperses, or spreads out. The more it spreads out, the weaker the signal becomes. This phenomenon is known as **free space path loss**; no action of the antenna can counteract it. The antenna is only a piece of equipment that radiates or absorbs a signal. Antenna gain can occur in response to a change in the antenna's shape by focusing RF radiation into a tighter beam of energy, which is measured by beam width in horizontal and vertical degrees. For example, focusing a 360-degree beam into a directed 30-degree beam causes the signal to travel much farther without requiring additional power.

There are three basic types of antennas: omnidirectional (also known as dipole), semidirectional, and highly directional. Table 6-4 lists their characteristics.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Type of antenna	Radiation direction	WLAN installation
Omnidirectional	All directions equally	Most common for a WLAN AP in one building
Semidirectional	Focuses energy in one direction	Used for short- or medium-range remote wireless bridges, such as a point-to-point link between neighboring buildings
Highly directional	Tightly focuses signal beam in one direction	Long-distance, point-to-point links, as in a wireless link between buildings separated by up to 25 miles; not commonly used for WLANs

© Cengage Learning 2014

Table 6-4 Basic antenna types

Remote Wireless Bridges

A remote wireless bridge connects wired and wireless network segments, much as an AP does, with two major exceptions. First, a remote wireless bridge transmits at higher power than an AP, which increases its range. Second, instead of the omnidirectional transmission used by APs, a remote wireless bridge uses a directional antenna to focus transmissions in one direction. You can use remote wireless bridges to connect wireless segments in a point-to-point configuration, as shown in Figure 6-12, or in a point-to-multipoint configuration to connect network segments, as shown in Figure 6-13.

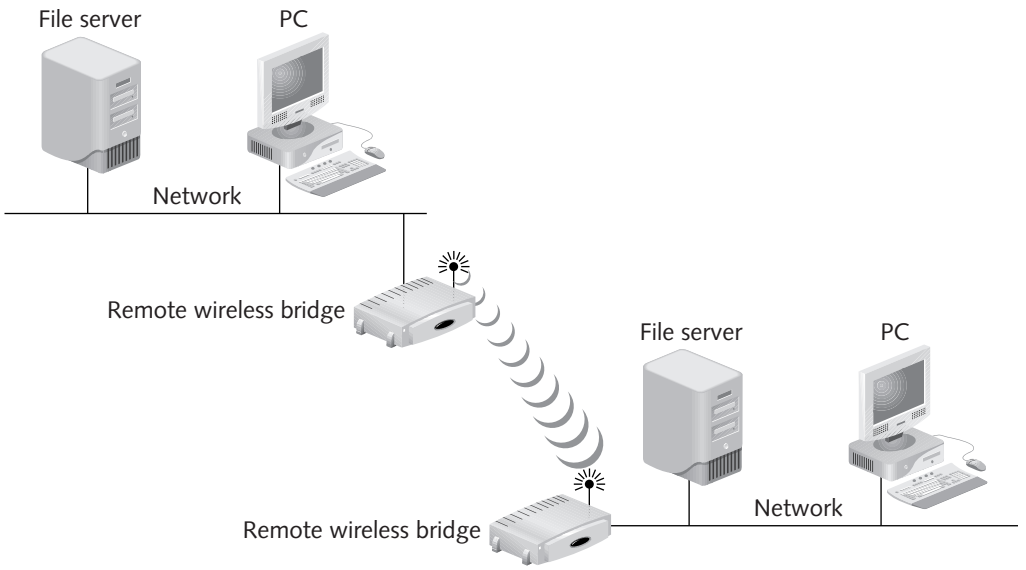

© Cengage Learning 2014

Figure 6-12 Point-to-point wireless bridging

A remote wireless bridge operates in four modes:

- *Access point mode*—The bridge functions as a standard AP and does not communicate with other remote wireless bridges.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

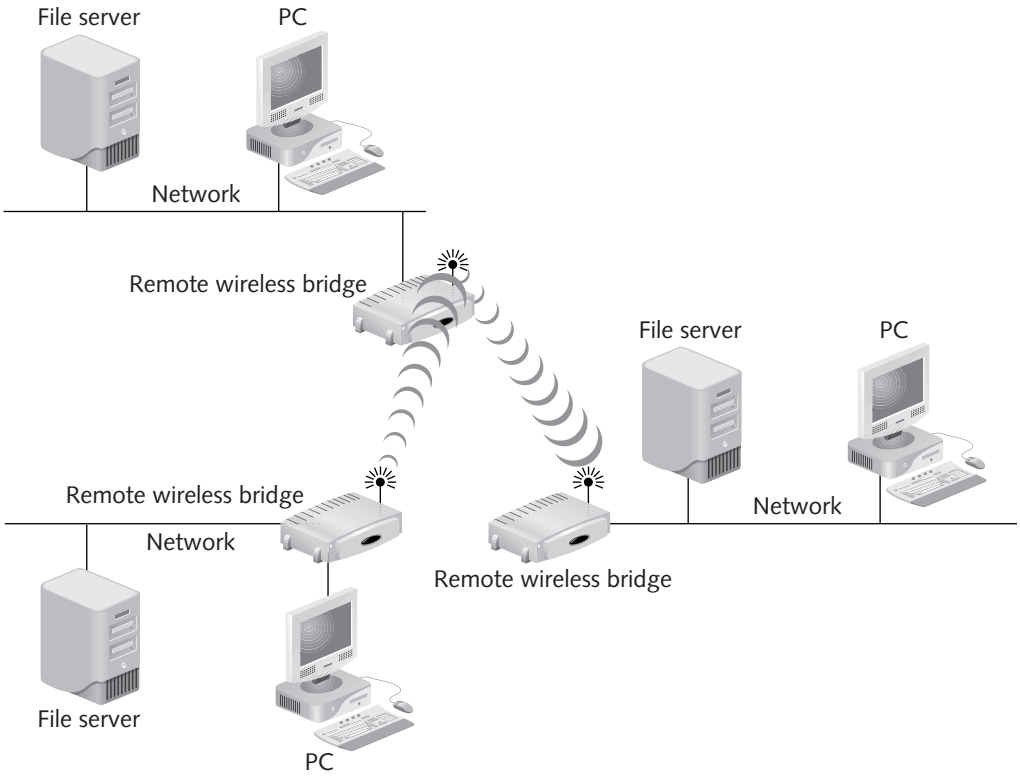

© Cengage Learning 2014

Figure 6-13 Point-to-multipoint wireless bridging

- *Root mode*—The bridge, called the root bridge, can communicate only with other bridges that are not in root mode. It cannot communicate with other root bridges or wireless clients.
- *Nonroot mode*—The bridge can communicate only with root bridges. Some manufacturers make remote wireless bridges that can also be configured as APs, enabling them to communicate with wireless clients and the root bridge simultaneously.
- *Repeater mode*—The bridge extends the range, as a hub or repeater does in a wired network. Repeater bridges can be used to extend the distance between network segments and can be positioned between other bridges.

Wireless Gateways

A wireless gateway combines management and security into a single appliance and can perform the following functions:

- Authentication
- Encryption
- Intrusion detection

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Malicious program protection
- Bandwidth management
- Centralized network management

WLAN Configurations

The Wireless LAN standards developed by the Institute of Electrical and Electronics Engineers (IEEE) define three basic WLAN configurations: Basic Service Set, Extended Service Set, and Independent Basic Service Set. In a **Basic Service Set (BSS)**, a group of wireless devices is served by a single AP. A BSS is sometimes called infrastructure mode. A BSS must be assigned a unique identifier known as the **service set identifier (SSID)**, which is the network (domain) name for the BSS. The geographical coverage for the BSS is called the **Basic Service Area (BSA)**, and the actual coverage varies depending on environmental factors, types of equipment, and number of users. Some devices are advertised as having a range of up to 375 feet from the AP, but their range is usually much shorter. Figure 6-14 shows a BSS configuration.

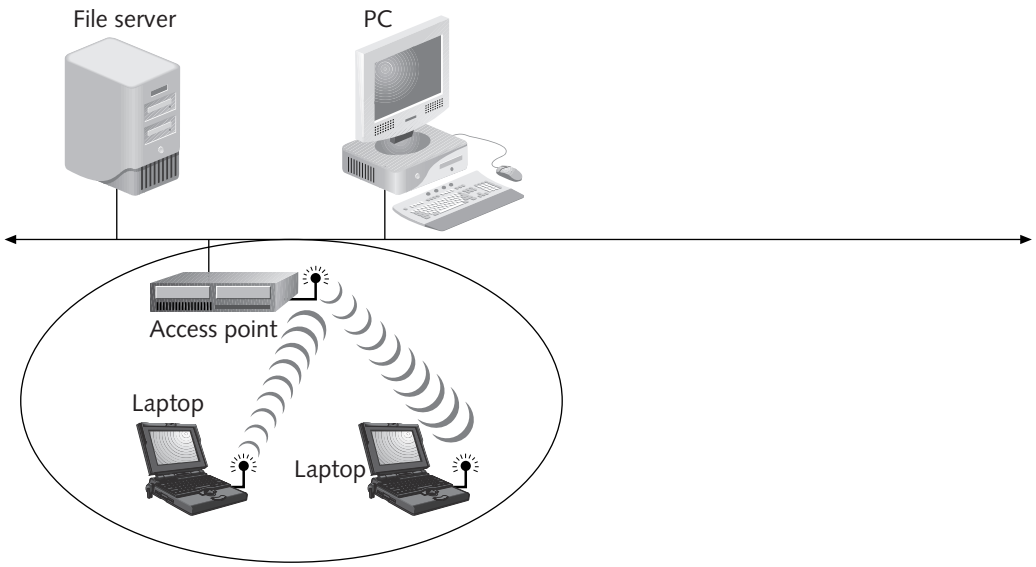

© Cengage Learning 2014

Figure 6-14 BSS configuration

For an **Extended Service Set (ESS)**, BSS networks are connected through a common distribution system, and APs are strategically placed so that users can move through a wider BSA. Although APs are often configured with different channels to prevent co-channel interference, the range of channels (the WLAN's frequency range) is part of the ESS. Like a BSS, the entire ESS network has a unique SSID.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

In an ESS (see Figure 6-15), APs are set up to provide some overlap. The coverage areas are called cells, and the WLAN is set up like a cell phone system. Movement between AP coverage areas is called **roaming**. When mobile users move into a new coverage area, their devices select the AP with the best signal strength. After the AP accepts the roaming device, the device tunes to the AP's channel but continues to survey radio frequencies at regular intervals to see whether another AP can provide a stronger signal.

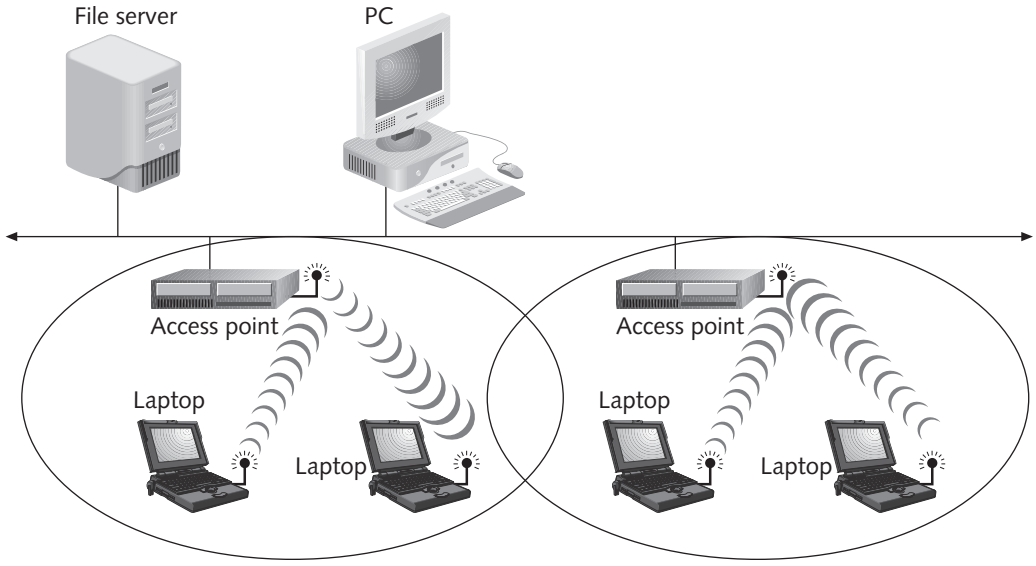

© Cengage Learning 2014

Figure 6-15 ESS configuration

Finally, an **Independent Basic Service Set (IBSS)** is a wireless network that does not use an AP. Instead, wireless devices communicate directly with one another, like computers in a peer-to-peer network. In wireless, IBSS is also known as ad-hoc mode.

Wireless Networking Standards

Like most computing technologies, wireless networking technology was developed in a haphazard way. Different companies and standards organizations worked on the same problems at the same time and came up with different solutions in many cases. Ultimately, the wireless standards process became more efficient, but still there are overlaps and uncertainty as wireless networking expands and its technology advances. In this section, you explore the important wireless networking standards.

IEEE 802.11

The IEEE created the 802 standard for networking. The number refers to the fact that the initial work was completed in February (the second month) of 1980. Standards like 802.2

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

and 802.3 define Ethernet, and the 802.5 standard defines token ring networks. The 802.11 standard for wireless networking was first released in 1997. The most recent iteration is called **IEEE Std. 802.11-2007**, and it includes all the ongoing amendments up to that time. Since 2007, other important amendments such as 802.11n (2009) have been added. In this section, you learn about the significant 802.11 standards.

802.11b (1999) One would think that 802.11a would have been the first wireless standard, but 802.11b was actually ratified before the 802.11a standard. Defined in Clause 18 of IEEE Std. 802.11-2007, this standard operates in the 2.4 GHz band. The maximum bandwidth supported is 11 Mbps. This standard is no longer used in contemporary WLANs.

802.11a (1999) Ratified later in the same year as 802.11b, the 802.11a standard now appears as Clause 17 of IEEE Std. 802.11-2007. This standard operates in the 5 GHz band and thus is not subject to interference by devices like microwave ovens and cordless telephones that operate at the 2.4 GHz range. The maximum bandwidth is 54 Mbps. Although 802.11a is not backward compatible with less expensive 802.11b/g devices, 802.11a manufacturers have developed devices that support a, b, and g standards.

802.11g (2003) Now described in Clause 19 of IEEE Std. 802.11-2007, 802.11g operates in the 2.4 GHz band, as does 802.11b, making the two standards interoperable. 802.11g also matches the 802.11a standard in its 54 Mbps bandwidth, but it shares 802.11b's problems with interference caused by devices operating in the 2.4 GHz band.

802.11i (2004) While manufacturers and customers focused on increasing wireless bandwidth, there was not much interest in wireless security. Wired Equivalent Privacy (WEP), as you learned in Chapter 5, was the only wireless security protocol that provided encryption. Once WEP was cracked, the industry rushed to develop another security protocol for wireless networking. The IEEE was deliberate in its approach; in the meantime, the Wi-Fi Alliance, formed by representatives of a number of wireless product vendors, developed Wi-Fi Protected Access (WPA). Once 802.11i was ratified, the WPA2 was released to map exactly to the 802.11i standard (Clause 8 of IEEE Std. 802.11-2007). You learn more about these protocols in Chapter 7.

802.11r (2008) Voice over Internet Protocol (VoIP) has become a popular method of reducing business costs by using the same network for data and telephone communications, which eliminates the need to use the phone company. As Voice over Wi-Fi (VoWiFi) becomes more important, wireless providers will have to address latency issues associated with roaming from one AP to another. 802.11i, otherwise known as WPA2, is sufficiently secure in an enterprise environment, but it takes too long to perform authentication—700 milliseconds or more. This delay only adds to the problems of implementing real-time applications like VoWiFi in enterprise networks. 802.11r is designed to provide **fast basic service set transition** (**FT**), which involves having a client perform a security association with the next AP before the client leaves the range of the current AP.

802.11n (2009) This amendment defines a standard that supports **multiple-input multiple-output** (**MIMO**). By using both 2.4 GHz and 5 GHz radio frequencies to simultaneously send

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

or receive data, bandwidth and range are both increased. Most implementations in business environments today include mixed standards of 802.11 a/g/n, which decreases the maximum supported bandwidth. However, bandwidth can reach 450 Mbps in all 802.11n networks.

802.11v (2011) This amendment, *Wireless Network Management*, defines standards that allow wireless stations to exchange operational information to improve wireless network performance. Wireless stations share information regarding network state, channel usage, diagnostics, interference, location, and topology.

802.11ac (Draft) This standard is expected to be ratified in late 2013. The first 802.11ac-compatible chipset was produced by Quantenna in late 2011. This standard will use the 5 GHz band, and it is expected to provide multistation WLANs with a bandwidth of 1 Gbps. The increased bandwidth will be produced by increased MIMO streams, enhanced modulation techniques, and wider-bandwidth channels.

Radio Frequency and the FCC

Because wireless primarily uses RF, it can interfere with radio and TV broadcasts, military and aviation communications, maritime communications, emergency services communications, satellite navigational systems, radar, cellular telephone services, and other critical applications. For this reason, most radio and some infrared and microwave frequencies are regulated strictly by the **Federal Communications Commission** (**FCC**) in the United States and by similar agencies internationally.

The FCC regulates what frequencies wireless communications can use, how much power antennas can emit, and other matters concerning the use of radio waves, infrared, and microwaves for communication. If your transmission strengths, frequencies, and directions violate FCC rules, your organization might face heavy fines. If your transmission interferes with police or fire service bands, your company could be held liable for injuries or deaths caused by that interference. When planning wireless deployment, be sure to check with your local FCC office to learn about regulations or requirements you must meet. Many frequencies are licensed to commercial, industrial, government, and military entities, and can be used only by the licensee. Even though wireless hardware vendors must sell devices that meet FCC regulations, some locations might have special circumstances.

Now that you understand the basics of how radio waves work and how they are used to form wireless networks, you can learn about the specifics of wireless security in the next chapter.

Chapter Summary

- Wireless transmissions use electromagnetic (EM) radiation, specifically radio frequency (RF) waves or infrared (IR) radiation, to communicate instead of sending electrical signals over physical cabling. RF is used most commonly in wireless networks. Infrared is mainly used for communication with peripheral devices.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- EM radiation travels in waves. A wave's maximum departure from the undisturbed state is the amplitude. Frequency is the number of times an event occurs in a specified period, and wavelength is the distance between repeating units (usually the midpoint) of the wave.
- The RF spectrum is divided into bands based on frequency. Most WLANs use frequencies in the 2.4 GHz and 5 GHz bands.
- The speed and transmission range of a wireless network vary depending on the standard, equipment, environmental factors, number of users, location of clients, and purpose.
- RF is highly susceptible to interference from solar activity, electrical storms, man-made radiation, and other RF emissions. Co-channel interference, multipath distortion, and fading are common types of interference.
- RF signal behavior is characterized by whether a factor contributes to a gain or loss of power. RF power is measured on a linear scale using milliwatts (mW) or on a relative scale using decibels (dB). dBm is the reference point that relates the decibel scale to the milliwatt scale.
- RF transmits a carrier signal, which can be changed via modulation to allow RF radiation to carry information. You can change a carrier signal's amplitude, frequency, or relative starting point.
- RF data can be analog or digital. Analog modulation methods are amplitude modulation, frequency modulation, and phase modulation. Digital modulation methods include amplitude shift keying, frequency shift keying, phase shift keying, and frequency division multiplexing.
- Spread spectrum spreads a narrowband signal over a broader portion of the RF band. The major methods are direct sequence spread spectrum, frequency hopping spread spectrum, and orthogonal frequency division multiplexing.
- Wireless network components include wireless NICs, access points, antennas, remote wireless bridges, and wireless gateways.
- Antennas transmit and receive radio waves, and can be omnidirectional, semidirectional, or highly directional.
- A remote wireless bridge operates in four modes: access point, root, nonroot, and repeater. Wireless gateways act as centralized management and security appliances, and can perform authentication, encryption, intrusion detection, malicious program protection, bandwidth management, and centralized network management.
- IEEE 802.11 standards define three WLAN configurations: Basic Service Set (BSS), Extended Service Set (ESS), and Independent Basic Service Set (IBSS).
- Wireless communication has several standards. The IEEE 802.11 Wireless Ethernet standards are used most often for WLANs in the United States. The 802.11 standards include 802.11a, 802.11g, and 802.11n.
- RF is subject to strict regulations by the FCC because of the potential for interference with critical communications, including radio, TV, military, and emergency services.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Key Terms

absorption A type of loss in which certain materials absorb the RF signal.

access point (AP) A wireless device that acts as the base station for a wireless network segment and serves as the link between wired and wireless segments.

amplitude The maximum departure of a wave from its undisturbed state; the height of the carrier wave.

amplitude modulation (AM) An analog modulation method in which the height of the carrier wave is changed.

amplitude shift keying (ASK) A binary modulation technique similar to AM that changes the height of the carrier wave. With ASK, the 1 and 0 bits do not both have carrier signals; the 1 bit has positive voltage and the 0 bit has zero voltage.

analog RF signal A continuous RF wave that oscillates between positive and negative voltage.

antenna A device used to transmit or receive radio waves. One end of an antenna usually points into the air and the other end is fixed to the ground directly or via a grounded device. Antennas are generally made of copper wire or other conductive material.

Basic Service Area (BSA) The geographical area of coverage for a WLAN.

Basic Service Set (BSS) An IEEE-defined WLAN configuration in which a group of wireless devices is served by a single AP.

carrier signal The signal on which a radio transmission is sent.

chipping code The bit pattern used in direct sequence spread spectrum. See also *direct sequence spread spectrum (DSSS)*.

co-channel interference Interference from two or more wireless access points transmitting with the same frequency.

dBm (decibel-milliwatt) A measure of relative signal strength; 1 mW = 0 dBm.

decibel The unit of measurement of the difference in strength between two signals.

diffraction A phenomenon that is similar to refraction, except that the signal bends around objects in its path; usually, these objects have a rough surface. See also *refraction*.

diffused IR transmission An infrared transmission that relies on reflected light rather than a line-of-sight connection between the IR emitter and detector.

digital RF signals RF signals that oscillate between positive and negative voltage but are broken into discrete segments with many starts and stops in the data stream. By contrast, analog transmissions have a constant stream of data.

direct sequence spread spectrum (DSSS) A wireless modulation technique that uses an expanded redundant code to transmit each data bit. The spreading code (key) is the chip sequence used to represent message bits, and the key is applied at the data level.

directed IR transmission A form of IR transmission that requires the emitter and the detector to be pointed directly at one another.

electromagnetic (EM) radiation Electromagnetic energy traveling as a self-propagating wave and spreading out as it moves.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

EM spectrum All forms of EM radiation collectively. From highest energy to lowest, EM radiation forms are gamma rays, x-rays, ultraviolet light, infrared light, visible light, microwaves, and radio waves.

Equivalent Isotropically Radiated Power (EIRP) The power radiated by a wireless system's antenna, including the intended transmission power and the unintended gain.

Extended Service Set (ESS) An IEEE-defined WLAN configuration in which multiple BSS networks are connected through a common distribution system.

fading A form of interference/power loss in which the signal is reflected or refracted by the atmosphere, ground, or other obstacles, and the reflected path interferes with the direct path.

fast basic service set transition (FT) A method specified in 802.11r that allows decreased latency in the authentication of a wireless client with an AP to which it is roaming.

Federal Communications Commission (FCC) The primary regulatory agency in the United States for wireless communication.

free space path loss The loss of signal strength resulting from the dispersion of the signal over distance.

frequency The number of times an event occurs in a specified time period; the rate at which an electrical current alternates, creating different radio transmissions.

frequency division multiplexing (FDM) An RF signaling method that enables multiple users to share a communication channel. Base signals are modulated on different carrier waves and combined to form a composite signal.

frequency hopping spread spectrum (FHSS) A modulation technique in which the range of frequencies used for transmission changes during the transmission. The time spent on a particular frequency is called the dwell time.

frequency modulation (FM) An analog modulation technique that changes the number of waves used to represent one cycle, so that the number of waves representing a 1 bit is greater than the number of waves representing a 0 bit.

frequency shift keying (FSK) A binary modulation technique that changes the frequency of the carrier signal.

Fresnel zone An area in which you can determine RF signal strength, identify potential obstacles, and determine the potential for multipath distortion between sending and receiving antennas.

gain The positive difference in amplitude between signals.

hopping code The sequence of changing frequencies in FHSS that determines the sequence and speed of frequency hops.

IEEE Std. 802.11-2007 A series of wireless networking standards that specify transmission speeds, ranges, and other aspects of wireless networks.

Independent Basic Service Set (IBSS) A wireless network that does not use an AP; also known as ad-hoc mode.

infrared (IR) radiation A type of radiation with a shorter wavelength than radio waves but a longer wavelength than visible light.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

infrared transmissions Transmissions that use infrared light pulses and that require an emitter and a detector. The emitter is usually a laser diode or LED. The detector is sometimes combined with the emitter and does just what its name implies: It detects IR transmissions and generates the corresponding electrical current for wired communication.

intentional radiator A device designed to generate radio signals, not including those from the antenna.

loss The negative difference in amplitude between signals.

modulation The process of varying the wave form to transmit a signal, such as modulating between on and off to create 1s and 0s.

multipath A signal that has more than one path from the transmitter to the receiver.

multiple-input multiple-output (MIMO) The use of multiple antennas operating at 2.4 GHz and 5 GHz, which allows for increased bandwidth in the 802.11n standard.

narrowband An RF transmission on one frequency or a very small frequency range.

orthogonal frequency division multiplexing (OFDM) A modulation technique that breaks a high-speed signal into smaller pieces and sends it in parallel across many lower-speed channels.

phase modulation (PM) An analog modulation technique that changes the cycle's relative starting point when the bit being transmitted changes from 1 to 0.

phase shift keying (PSK) A binary modulation technique that starts and stops the signal to represent a binary digit.

polarization The plane on which radio waves propagate, or the orientation of the radio waves as they leave the antenna; usually the orientation is vertical or horizontal.

Power over Ethernet (PoE) A technology that sends power for the AP unit over unused wire pairs in standard twisted-pair Ethernet cabling.

radio frequency (RF) waves A type of EM radiation used most commonly for wireless networks.

reflection A type of loss in which the RF wave bounces off certain materials.

refraction A phenomenon in which signals traveling through different air masses can be bent rather than traveling in a straight line. The signal changes in response to atmospheric conditions.

roaming A wireless device moving between access points, or cells, without losing its ability to transmit and receive. Cells are interconnected so that as the wireless device moves between cells, it negotiates a new connection depending on signal strength.

scattering A type of loss in which small objects, such as raindrops, dust, or foliage, scatter the signal.

service set identifier (SSID) A unique identifier of up to 32 characters that serves as the wireless network name.

spread spectrum An alternative to narrowband transmissions that spreads a narrow signal over a broader portion of the radio band.

voltage standing wave ratio (VSWR) A type of loss caused by differences in the equipment itself. For example, if one part of the equipment has a different impedance than others, the

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

signal can be reflected back within the device, and the reflected power can cause the device to burn out.

watt A measure of power or the rate at which work is done.

wave A means of transporting energy from one place to another without physical movement of material.

wave propagation The pattern of wave dispersal as waves travel from sending to receiving antennas.

wavelength The distance between two successive amplitude peaks.

wireless networking Any exchange of data between computers and other devices that uses RF transmission methods instead of traditional wired cables.

Review Questions

1. In which of the following situations would a wireless network be an effective alternative?
   a. a business that occupies temporary space
   b. a network with employees who travel
   c. an older building with no wiring
   d. all of the above
2. Wireless communication emits which types of EM radiation? (Choose all that apply.)
   a. gamma rays
   b. infrared radiation
   c. ultraviolet light
   d. radio frequency waves
3. The maximum departure of a wave from its undisturbed state is called which of the following?
   a. frequency
   b. wavelength
   c. amplitude
   d. hertz
4. All types of electromagnetic radiation are collectively called which of the following?
   a. EM spectrum
   b. EM field
   c. light spectrum
   d. visible light spectrum

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

5. Frequency is defined as which of the following?
   a. the distance between repeating units of a wave
   b. the number of times an event occurs in a specified period
   c. a method of transporting energy without physical movement of material
   d. radiation spreading out as it moves
6. The distance between midpoints of a wave is called which of the following?
   a. frequency
   b. wavelength
   c. amplitude
   d. hertz
7. What type of infrared transmission relies on reflected light?
8. Infrared wireless is extremely fast, can travel long distances, and is not susceptible to interference. True or False?
9. Which of the following is a common type of signal loss? (Choose all that apply.)
   a. deletion
   b. diffraction
   c. deflection
   d. scattering
10. Refraction occurs when the RF signal is dispersed by small objects, such as raindrops or foliage. True or False?
11. Which of the following is used to provide a relative measurement of RF power?
    a. milliwatts
    b. decibels
    c. watts
    d. 10s and 3s rules of RF math
12. Why is digital modulation superior to analog modulation? (Choose all that apply.)
    a. It makes more efficient use of bandwidth.
    b. It has fewer interference problems.
    c. Error correction is more compatible with other digital systems.
    d. Less power is required to transmit.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

13. Which of the following is considered an advantage of spread spectrum over narrowband wireless transmission? (Choose all that apply.)
    a. increased security
    b. decreased susceptibility to interference
    c. requires less power to reach the same amplitude
    d. does not require a chipping code
14. Radio waves striking an antenna create infrared radiation. True or False?
15. IEEE 802.11ac is expected to support a bandwidth of _____.
    a. 54 Mbps
    b. 102 Mbps
    c. 512 Mbps
    d. 1000 Mbps

Hands-On Projects

Hands-On Project 6-1: Researching the FCC

Time Required: 30 minutes

Objective: Explore the FCC Web site to learn more about wireless communication.

Description: In this project, you explore the FCC Web site to learn about frequencies that are commonly used for wireless transmission. You also use the FCC license search tool to see what types of wireless communication entities operate in your area.

1. Start your Web browser and go to **http://reboot.fcc.gov/reform/systems/spectrum-dashboard**.
2. Click **Browse Spectrum Bands**. In the Select Using Frequency Band window, click the **+** button on the Zoom slider bar twice to enlarge the spectrum.
3. Hold your cursor over the right arrow on the right side of the spectrum until the 2450–2483.5 MHz band is within the viewable range of the spectrum. Hold your cursor over the range to see a description. (See Figure 6-16.)
4. What 802.11 standards operate in this range?

   _____
5. Click the **Back** button to return to the Spectrum Dashboard main page. Click **Browse Using a Map**. Click a state to populate the State and Counties windows below the map. Select a county and click **Search**. View the results to see what RF bands are licensed in the county.
6. Go to **http://reboot.fcc.gov/license-view/**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

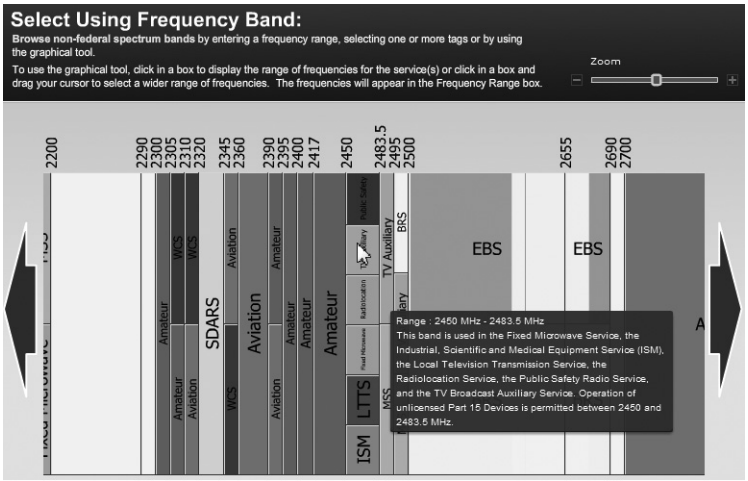

Source: FCC

Figure 6-16 FCC spectrum dashboard

7. What organization holds the most active licenses?

_____

8. Leave your system running for the next project.

Hands-On Project 6-2: Installing inSSIDer

Time Required: 15 minutes

Objective: Download and install a freeware program for detecting wireless networks.

Description: In this project, you install inSSIDer, a free wireless networking utility for Windows. If you do not have a WNIC and wireless AP, this project at least familiarizes you with the interface and some features of the program.

1. Start your Web browser and go to **http://download.cnet.com/inSSIDer/3000-18508_4-10848357.html?tag=mncol;4**.
2. Click **Download Now**.
3. When the File Download window opens, click **Save**. Click **View downloads**, and click **Run** to the right of the inSSIDer row. If the User Account Control window opens, click **Yes**.
4. In the CNET Download.com window, click **Next Step**. If a Special Offer window appears, click the **I do not accept** option button, and click **Next Step**. Click **Install Now**.
5. In the Welcome to the inSSIDer Setup Wizard window, click **Next**. In the Select Installation Folder window, click **Next**. In the Confirm Installation window, click **Next**. In the Installation Complete window, click **Close**.
6. Click **Start**, click **All Programs**, click **MetaGeek**, and click **inSSIDer**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

7. In the upper-right corner of the inSSIDer window, click **Start** if necessary. Click the **Time Graph** tab. If the program has detected that a supported WNIC and APs are present, your screen will look similar to Figure 6-17.

Source: inSSIDer

Figure 6-17 inSSIDer main window

8. Examine the information. The most important columns are MAC Address, which shows the Data Link layer address of the WNIC; SSID, which shows the WLAN name; and RSSI, which shows signal strength. Although the RSSI numbers are negative, note that the higher the number, the stronger the signal. For example, 30 dBm is a very strong signal, and -90 dBm is a very weak signal. Other columns include Channel and Security. Compare the RSSI column with the time graph, which displays changes in signal strength over time.
9. In Figure 6-17, three SSIDs have the same name: Far. How can you explain this?

_____

_____

_____

As a clue to answering the question in Step 9, examine Figure 6-18, which shows the 2.4 GHz Channels tab, and Figure 6-19, which shows the 5 GHz Channels tab.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

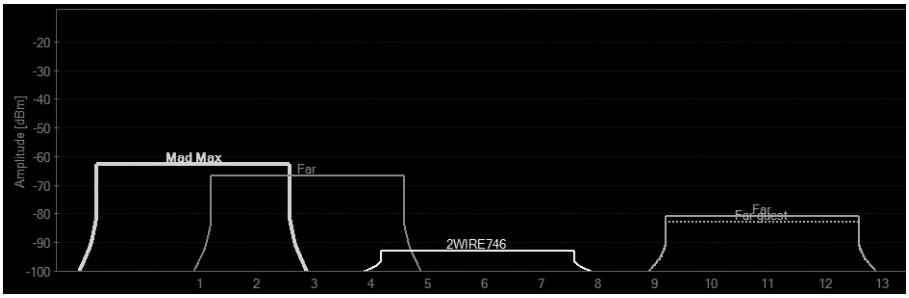

Source: inSSIDer

Figure 6-18 inSSIDer 2.4 GHz Channels tab

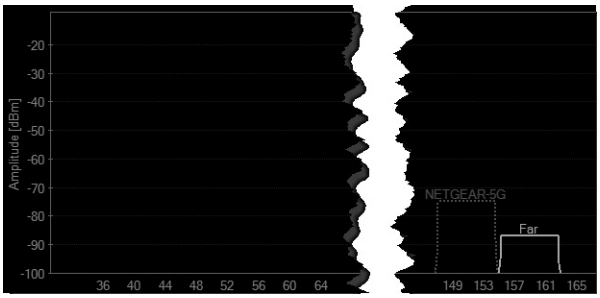

Source: inSSIDer

Figure 6-19 inSSIDer 5 GHz Channels tab

10. List some ways that a WLAN administrator could use inSSIDer.

11. Leave your system running for the next project.

Hands-On Project 6-3: Researching Access Points

Time Required: 30 minutes

Objective: Research the Internet to identify several wireless access points.

Description: In this project, you research wireless access points from different vendors. You determine which APs are suitable for home use and which are suitable for businesses. You also examine at least one enterprise-class access point.

1. Start Microsoft Word or another word-processing program, and create a table with four columns and five rows. Label the columns Access Point, Features, Cost, and Recommended Use.
2. Start your Web browser and search for wireless access points.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

3. Visit vendor Web sites to gather information on at least five APs, including one enterprise-class AP. Enter the information in the table you created in Step 1.
4. Turn in the finished table to your instructor.

Hands-On Project 6-4: Researching Wireless Signal Interference

Time Required: 15 minutes

Objective: Learn how common materials can cause passive interference with wireless signals.

Description: In this project, you examine a physical location, list the objects that might cause interference with wireless signals, and note how severe the interference could be. Remember that a wireless transmission can travel in all directions, so look up and down, too.

1. Start your Web browser, go to **http://support.gateway.com/s/Mobile/Gateway/450ROG/3501353faq108.shtml**, and read the information.
2. Select a point in a room where you might locate an AP. Choose locations for wireless clients as well.
3. Examine all areas in the room. Do not forget to consider what is behind the walls and above and below your AP.
4. Write down all items that could cause communication problems for a wireless network.
5. Turn in your finished paper to the instructor or use it for class discussion.

Case Projects

Case Project 6-1: Planning Fixed-Mobile Access

You work for a network consulting firm, and you have been assigned to prepare a presentation for the executives of an amusement park company who are concerned about communication costs and session persistence when employees use their company-issued cellular phones. The network at the park includes extensive 802.11-based APs inside buildings, but many outdoor locations at the park have no WLAN coverage. Employees often use cellular phones for voice communications because they are frequently moving around the park. To save costs, the executives are interested in having the cellular phones use the WLAN when it is in range.

The executives have relatively little training in network technology, so you need to present the concepts of *fixed-mobile convergence* and *internetworking with external networks* in a way that the executives can understand. Prepare a slide presentation and a 5- to 10-minute discussion to address the executives' concerns. You also need to address how to modify the WLAN to support fixed-mobile convergence.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 7

Understanding Wireless Security

After reading this chapter and completing the exercises, you will be able to:

- Explain basic concepts in wireless security
- Describe security solutions for wireless networking
- Describe security solutions for handheld wireless devices

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

This chapter introduces you to the fundamentals of wireless security. You learn about the threats that wireless networks face and solutions for defending against them. You learn about passive attacks, active attacks, rogue devices, and key management concerns so you can identify major weaknesses in wireless network design and determine possible solutions.

In most organizations, the people who use the systems are simultaneously the most critical resource and the biggest threat to security. Therefore, comprehensive security policies and security training programs are crucial to ensure the secure use of mobile devices. Countermeasures to security threats include physically securing devices, using strong authentication and key management, and using strong encryption to protect data stored on portable devices or transmitted wirelessly.

Security Concerns of Wireless Networking

Wireless networks face different threats than wired networks. Wireless networks do not have physical cabling to secure, so packets literally travel the airwaves and are vulnerable between the transmitter and receiver. If no encryption is used, everything is sent in cleartext, including passwords and confidential data.

This section explains the evolution of wireless security concerns. Wireless communication relies heavily on the Media Access Control (MAC) sublayer of the Data Link layer in the OSI model. MAC frames and MAC addresses play an important role in wireless communication, but they also create vulnerabilities.

Next, you learn about passive and active scanning of wireless networks. Scanning for wireless signals is a valid activity because it allows wireless stations to find and connect to available networks, but these methods are also used to find networks to attack.

The basic authentication methods in IEEE 802.11 networks also cause security problems, so you learn about the inherent vulnerabilities of IEEE 802.11's authentication mechanisms. Then you examine wireless network attacks and major security vulnerabilities, such as the challenge of managing keys and the dangers of using default settings.

Next, you examine common methods for securing wireless networks. Wi-Fi Protected Access (WPA), WPA2, and 802.1x are examples of robust solutions for securing modern wireless networks.

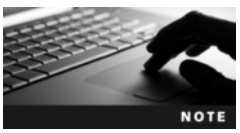

In wireless networks, a wireless device is called a *station*, which is similar to the term *node* in wired networks.

IEEE 802.11 Media Access Control: Frames

The MAC sublayer of the OSI model performs many critical functions in a wireless network:

- Discovering wireless access points, channels, and signal strengths
- Joining the wireless network, including authentication and association to the access point

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Transmitting data
- Maintaining the connection

Each access point (AP) has a 0- to 32-byte service set identifier (SSID) that essentially functions as the name of the network. SSIDs separate airwaves into segments so that wireless networks in the same physical area can operate independently of one another. SSIDs can also be mapped to virtual LANs, so some APs support multiple SSIDs.

When wireless stations communicate, they use MAC frames to locate wireless networks, establish and maintain the connection, and transmit data. All MAC frames contain a control field that identifies the 802.11 protocol version, frame types, and codes that specify wireless configurations. Frames also contain MAC source and destination addresses, a frame sequence number, and a frame check sequence (FCS) for error detection.

The 802.11 standard has three types of MAC frames: management frames, control frames, and data frames. **Management frames** establish and maintain communications. They are always sent in cleartext, and many contain SSIDs. Even link encryption, such as Wired Equivalent Privacy (WEP), does not encrypt management frames. The security problem with management frames is that anyone who intercepts one can discover the SSID and then have part of the information needed to access the network. Figure 7-1 shows the structure of a management frame.

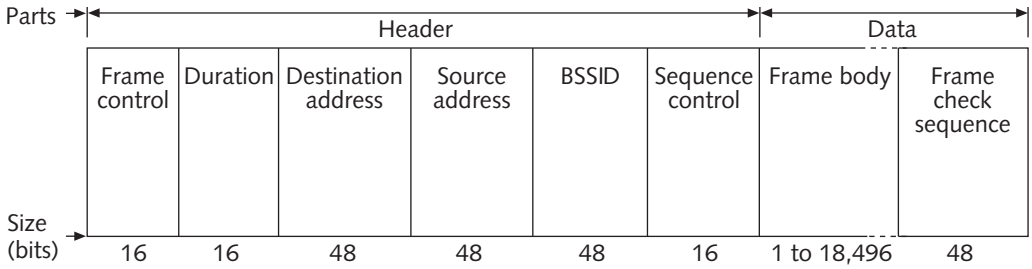

© Cengage Learning 2014

Figure 7-1 An IEEE 802.11 management frame

The following list explains the fields in a management frame:

- *Frame control*—Information such as the IEEE standard version and whether encryption is used
- *Duration*—The amount of time in microseconds needed for transmission
- *Destination and source address*—Source and destination addresses of sending and receiving stations
- *BSSID*—The basic service set identifier (the network name), a variation of SSID
- *Sequence control*—The packet's sequence number and fragment number
- *Frame body*—The data payload
- *Frame check sequence*—Error detection

Table 7-1 describes common types of management frames.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Frame type	Purpose
Association request	Allows an AP to allocate resources for a wireless station
Association response	Sent by the AP in response to an association request frame; indicates whether the request is accepted or rejected
Reassociation request	Sent to the new AP when a wireless station roams into a different AP coverage area (also called a "cell")
Reassociation response	Sent by the AP in response to a reassociation request frame; indicates whether the request is accepted or rejected
Probe request	Sent by a station when it needs information from another station
Probe response	Sent by a station in response to a probe request frame; indicates capabilities, supported data rates, and other information
Disassociation	Sent by a station to another station if the sender wants to terminate the connection
Authentication	Sent by the AP to determine whether to allow a wireless station to enter the network
Deauthentication	Sent by a station to another station if the sender wants to terminate the connection
Beacon	Sent by an AP to any listening stations to advertise services or information available on the wireless network; beacons contain SSIDs, capabilities, supported rates, and other information about the AP and wireless network

© Cengage Learning 2014

Table 7-1 Management frame types

Reassociation requests and responses might seem confusing because the station does not authenticate again. A station can be authenticated on multiple APs but can be associated with only one at a time. If a station leaves the network on which it is currently authenticated, it must authenticate to the new network. However, moving to a new coverage area does not necessarily mean that the station has changed networks; it simply means that the station has left one AP's coverage area and moved into another.

Control frames help deliver data frames between stations and control access to the medium. Figure 7-2 shows a typical control frame. There are six types of control frames; the following list explains the four most common types:

- *RTS*—A **request to send** (**RTS**) frame is the first step of the two-way handshake before sending a data frame. A station using the RTS/CTS mechanism sends an RTS frame when it wants to transmit data.
- *CTS*—In response to an RTS frame, a **clear to send** (**CTS**) frame gives a station clearance to send. It contains a time value that keeps all other stations from transmitting long enough to give the sending station time to transmit.
- *ACK*—After receiving a data frame, the receiving station performs error checking. If no error is found, the station sends an **acknowledgement** (**ACK**) frame; if no ACK frame is received, the sending station retransmits the frame.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *PS-Poll*—A **power-save poll** (**PS-Poll**) frame is used when a station has awakened from power-save mode and sees that an AP has frames buffered for it. The station sends this frame to let the AP know it can transmit the buffered frames.

Stations and APs have a configuration parameter called the RTS threshold that indicates whether the RTS/CTS process is used before transmitting. If a station is configured to use RTS/CTS frames, it transmits an RTS frame requesting access to the medium. After a CTS frame is received in response, the station knows it can transmit safely. All listening stations hear the exchange and hold transmissions long enough for data to be transmitted and to receive an ACK frame indicating success. This process is not foolproof, however. A station that is not listening or not connected to hear the exchange could still transmit during the reserved time.

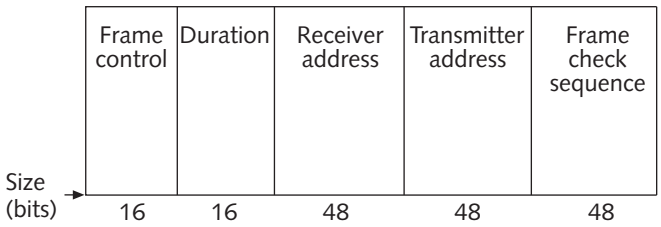

© Cengage Learning 2014

Figure 7-2 An IEEE 802.11 control frame

The following list explains the fields in a control frame:

- *Frame control*—Information such as the protocol version, frame type, and whether encryption is used
- *Duration*—The amount of time in microseconds needed for transmission
- *Receiver address*—The MAC address of the receiving station
- *Transmitter address*—The MAC address of the transmitting station
- *Frame check sequence*—Error detection

Data frames carry the TCP/IP datagram, and the payload is encrypted (see Figure 7-3). Data from higher-layer applications, such as printer control data or Web pages, is carried in the data frame body. Data frame fields labeled Address 1, Address 2, Address 3, and Address 4 carry the BSSID, source MAC address, destination MAC address, and address of the transmitter or receiver, respectively.

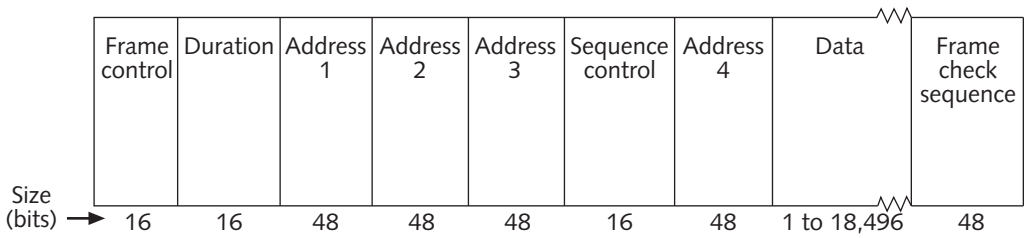

© Cengage Learning 2014

Figure 7-3 An IEEE 802.11 data frame

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The following list explains the fields in a data frame:

- *Frame control*—Information such as the IEEE standard version and whether encryption is used
- *Duration*—The amount of time in microseconds needed for transmission
- *Destination and source address*—Source and destination addresses of sending and receiving stations
- *BSSID*—The basic service set identifier (the network name), a variation of SSID
- *Sequence control*—The packet's sequence number and fragment number
- *Frame body*—The data payload (information from higher OSI layers)
- *Frame check sequence*—Error detection

Unlike MAC addresses and fully qualified domain names (FQDNs), SSIDs are not registered, so two wireless networks could use the same SSID. A station could also have a null SSID that allows it to match all SSIDs. If a beacon frame contains a null SSID, attackers just have to capture frames that contain the correct SSID. You can turn off beaconing for most current APs, but this measure is not very effective; attackers can wait for management or control frames that contain the information they want, or they can spoof (impersonate) management frames and sniff the responses to find information. **Sniffing** is capturing network traffic during transmission.

In addition, several management frames contain the network's SSID, and management frames are always transmitted in cleartext, even when encryption methods are used. A passive scan can reveal SSIDs to attackers easily, and some APs send beacon frames as often as several times a second, so attackers do not have to wait long to intercept frames that contain SSIDs.

Scanning and Attacks

When a wireless station wants to connect to a wireless network, it begins listening on each available channel for an AP's beacon frame broadcast. This listening process is called passive or active scanning.

In **passive scanning**, a wireless network interface card (WNIC) listens to each channel for a few packets, and then moves to another channel. Because the station listens without transmitting, its presence is not usually revealed. A WNIC's **radio frequency (RF) monitor mode** allows passive scanning, although many WNICs have this capability disabled by default in their firmware. In RF monitor mode, the WNIC's equivalent of promiscuous mode in NICs, a WNIC can capture packets without authenticating or associating with an AP or ad-hoc (peer-to-peer) wireless network. For this reason, passive scanning is difficult or impossible to detect.

A **passive attack** uses passive scanning to gather information about a wireless network for later use; this information includes SSIDs, MAC addresses, passwords, and usernames. Wireless networks are particularly vulnerable to passive attacks, such as sniffing and network reconnaissance. Because transmissions in wireless networks travel over airwaves, attackers simply need to be within range of an unsecured network to intercept packets and then analyze them to get more information for further attacks.

In **active scanning**, the station sends a probe request frame on each available channel and waits for a probe response frame from available APs. Even after a wireless station is associated with a

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

wireless network, it continues to scan for beacon frames. If a station is disconnected from the network, it can reconnect more quickly if it already has the information needed to connect to another AP or station. In general, stations select the strongest signal unless they are configured to connect to a specific AP or other criteria are set. Occasionally, a station connects to the first AP from which it receives a signal.

In **active attacks**, attackers use several techniques to probe wireless networks in an attempt to gather information. Unlike passive attacks, most active attacks can be detected by network security measures, although some bypass security measures and APs, sending traffic directly to the victim station. Table 7-2 lists some common active attacks; notice that some are denial of service (DoS) attacks.

Many free and commercial tools are available to help intruders automate attacks. Common tools include Aircrack-ng, coWPAtty, WPA Cracker, and WepAttack. However, be careful when you test wardriving software; you could connect to a wireless network inadvertently. If you cause harm, you could face charges whether your actions were intentional or not. As for the legality of simply connecting to an open wireless network, the issue is unclear in many parts of the United States. Some states, such as Florida, have found such connections to be illegal; however, as of this writing, the issue has not been tested in court in most states.

Attack	Method
Jamming (DoS attack)	The attacker floods airwaves with noise to weaken the RF signal and cause the wireless network to stop functioning.
Association flood (DoS attack)	The attacker authenticates several fake stations to send a flood of spoofed association requests, which overflow an AP's association request table. An AP can have up to 2007 concurrent associations before the association request table overflows and refuses further associations.
Forged disassociation (DoS attack)	The attacker sends a forged disassociation frame that contains the spoofed source MAC address of an AP. The target station attempts to reassociate, and the attacker continues to send disassociation frames to prevent reassociation or replies with a reassociation response.
Forged deauthentication	The attacker monitors transmissions to identify target stations. When a data frame or association response frame is captured, the attacker sends a spoofed deauthentication frame that contains the AP's MAC address. To prevent reconnection, the attacker continues to send deauthentication frames for the duration of the attack.
MAC address spoofing	The attacker inserts spoofed values in a frame's sender MAC address field.
Session hijacking	The attacker causes valid users to lose their connections by sending a forged deauthentication or disassociation frame to their stations, for example. The attacker then assumes their identities and privileges by sending their stations' MAC addresses to the AP. The attacker disables users' systems by using a DoS or buffer overflow attack so that their stations cannot reconnect.

© Cengage Learning 2014

Table 7-2 Common active attacks (*continues*)

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Attack	Method
Brute force	The attacker uses a program that attempts every possible key combination by changing one character at a time systematically; this attack is an attempt to decrypt a message and discover the default key.
Dictionary	In an attempt to determine passphrases, the attacker encodes dictionary words in the same way the passphrase was encoded. When the encoding matches, the attacker has found the passphrase.

© Cengage Learning 2014

Table 7-2 Common active attacks (*continued*)

Some of these attacks take advantage of Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA) being used as a wireless access method. With CSMA/CA, stations listen before transmitting. If traffic is already in transmission, stations wait until airwaves are clear before transmitting. In wireless DoS attacks, simply flooding airwaves with transmissions prevents legitimate stations from transmitting.

Wired Ethernet networks use the CSMA/Collision Detection (CD) access method. Stations listen for collisions and wait a random amount of time before transmitting when they detect a collision.

Wardriving and Exploitation of Rogue Devices

In **wardriving**, a potential attacker drives around with a laptop and WNIC in RF monitor mode to detect unsecured wireless signals. Wardrivers can use inexpensive hardware and free software to access an unsecured wireless network easily, regardless of whether it is part of an enterprise network or home network. A myth has developed that wardrivers are simply hijacking a connection, but in fact they are stealing information, violating data privacy, and possibly using their access to cause malicious damage.

Rogue devices are wireless devices that employees connect and use without authorization or verified configurations. Many administrators consider rogue devices to be a minor problem, but any unauthorized device represents a security vulnerability that can be exploited. In addition, rogue devices are usually configured poorly, so attackers can often locate them quickly and easily.

For example, Betty is an accounts payable manager who wants a mobile connection in the file room because she often reviews and files paperwork or reconciles accounts there. If she could take her laptop with her, she would not have to haul files between her office and the file room. She asked the IT Department to set up a wireless connection for her, but the request was refused. Undeterred, Betty researches wireless connections and finds out that she needs only a wireless router and WNIC. Her laptop has an onboard WNIC, so she buys an inexpensive wireless router. She follows the simple instructions to attach the router, and then uses the Windows 7 wizard to configure the WNIC and connect to the corporate network. She thinks that the IT Department was too lazy to help her because the wireless connection was not hard or time consuming to set up, and she is satisfied that she was technically savvy enough to do it herself.

The next morning, a wardriver is cruising around the neighborhood. He looks at his laptop and sees that his antenna has detected a WLAN broadcasting advertisements as "linksys."

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

He knows immediately that he has detected a wireless router configured with default settings. His open-source software tells him that the router is broadcasting unencrypted packets, so he connects to it and starts capturing packets. He examines the packets and finds what he is looking for: a valid username and password. With this information, he can probably gain full access to the network and steal sensitive information. With a little more effort, he can disable software, infect the network with a virus, and delete or corrupt data.

This situation happens every day, but companies often do not realize that well-intentioned employees are a common source of wireless security breaches. Worse, the IT Department might not have configured wireless network security and auditing correctly, so rogue devices often go unnoticed until a breach is discovered. According to a 2011 survey, 19% of wireless networks use the obsolete and unsecure WEP standard, 11% use default configurations, and 6% have no security measures at all. While these statistics show an improvement over similar surveys in the previous decade, the proliferation of wireless devices makes lax security measures a headache for security professionals and makes intrusions easy for attackers.

Although it might seem that Betty is solely responsible for any harm that occurs to the network, the IT Department shares responsibility. Too often, IT personnel are unresponsive or abrupt with users. Had an IT Department employee politely explained to Betty why her request could not be accommodated, she might not have taken steps to set up her own wireless connection. Unfortunately, she was not as knowledgeable as she thought she was, so her laptop's connection was not secured correctly. A comprehensive security-training program that included information about wardriving and rogue devices would have helped prevent Betty's actions. IT professionals should treat users with respect to help ensure that network security is maintained. If you work in IT and maintain a good working relationship with users, they are usually willing to follow security procedures. If you refuse to help them, they might circumvent the procedures you have set up.

Wireless Man-in-the-Middle Attacks

In a man-in-the-middle (MITM) attack, attackers intercept the transmissions of two communicating nodes without the users' knowledge. The transmissions can be modified and then forwarded to the intended destination, blocked from being delivered, or simply read and passed on. Most wired networks take countermeasures to reduce the risk of MITM attacks, but wireless networks provide new opportunities to use this attack method. A wireless MITM attack follows the same general procedure as on a wired network, but attackers often set up a fake AP to intercept transmissions and make stations think they are connecting to an authentic AP. Figure 7-4 shows a typical wireless MITM attack.

Secure WLAN Implementation

Association with a Wireless Network

To access services and resources, a station must be associated with an AP (in infrastructure mode) or another station (in ad-hoc mode). **Association** is a two-step process. First, a station listens for beacon frames to locate a network to join and then goes through the authentication process; a station cannot be associated without being authenticated first. Second, the station sends an association request frame. If the AP accepts, it reserves memory space for the

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

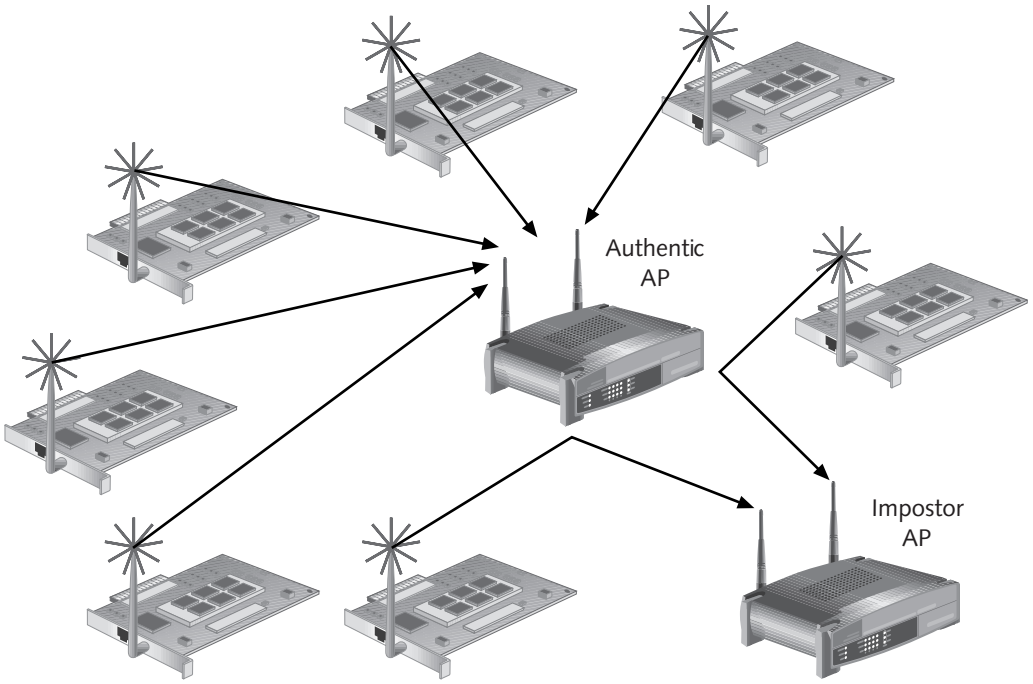

© Cengage Learning 2014

Figure 7-4 A wireless man-in-the-middle attack

station and sends back an association response frame that contains the association ID and connection information, such as supported data rates. A station can be authenticated on several APs simultaneously, but it can be associated with only one network at a time.

Wireless Authentication

When a wireless station wants to connect to a network, it authenticates to an AP or another station first in a process called **IEEE 802.11 authentication**. A key difference between wireless and wired networks is that the wireless station, not the user, is authenticated before being connected to the network. To access resources, users are then prompted to authenticate after the connection has been made. However, the AP typically does not perform user authentication; that process is handled by a Remote Authentication Dial-In User Service (RADIUS) or Terminal Access Controller Access Control System Plus (TACACS+) server, directory services, or another means of network authentication.

Early 802.11 standards provide two basic authentication methods: open system authentication and shared key authentication. In **open system authentication**, a station is authenticated without further checking as long as it has an SSID that matches the network it is attempting to join (see Figure 7-5). In ad-hoc mode, which is used in a peer-to-peer network, one station sends an authentication frame to another and receives a frame indicating recognition. In infrastructure mode—another term for a BSS configuration, with APs connecting stations—a station transmits its request to the AP and is authenticated as long as it has the correct SSID. Open system authentication provides little security because SSIDs are transmitted in manage-

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

ment frames in cleartext. Attackers can easily use passive scanning to find this information from other authenticated stations or from the AP if it is broadcasting beacon frames.

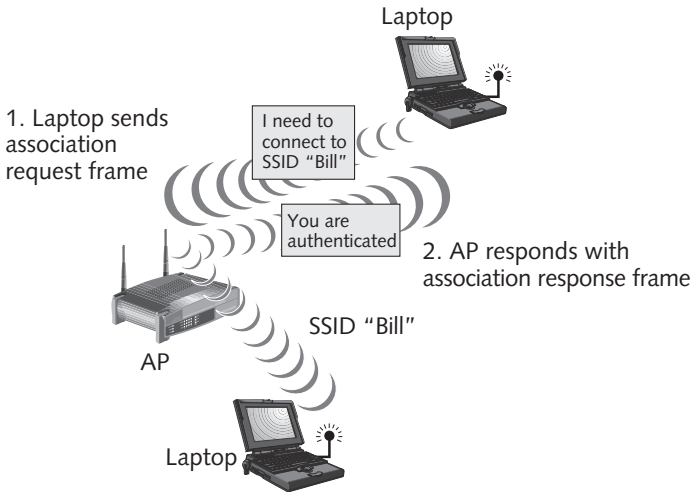

© Cengage Learning 2014

Figure 7-5 Open system authentication

Shared key authentication uses a standard challenge-response process with shared key encryption. A wireless station sends an authentication frame to an AP, which returns an authentication response frame that contains the challenge text. The station encrypts the text with its shared key and returns it to the AP. Using its own copy of the shared key, the AP decrypts the text and compares it to the original challenge text. If they match, the AP sends another authentication frame with the results, and the station is authenticated. If they do not match, the station's connection attempt is rejected.

A single shared key is distributed to all stations on a wireless network before they can communicate (see Figure 7-6). The IEEE standards do not specify any mechanisms for key management, so vendors, administrators, developers, and others must devise their own key management schemes. Managing shared keys effectively is a challenging aspect of securing wireless transmissions. Shared key authentication is also considered weak if it uses WEP for encryption. Attackers can use passive scanning to capture packets and crack the shared key; after they have the key, they can be authenticated to the network easily and then launch attacks.

The initial 802.11 standard included WEP and outlined several of the following cryptographic objectives, but did not specify how to achieve them:

- *Efficiency*—The encryption algorithm must be efficient enough to be used in hardware or software.
- *Exportable*—U.S. Department of Commerce security guidelines must be met so that WEP-enabled devices can be exported outside the United States.
- *Optional*—WEP must be an optional feature.
- *Self-synchronizing*—Each packet must be encrypted separately so that a single lost packet does not result in all packets being unreadable.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

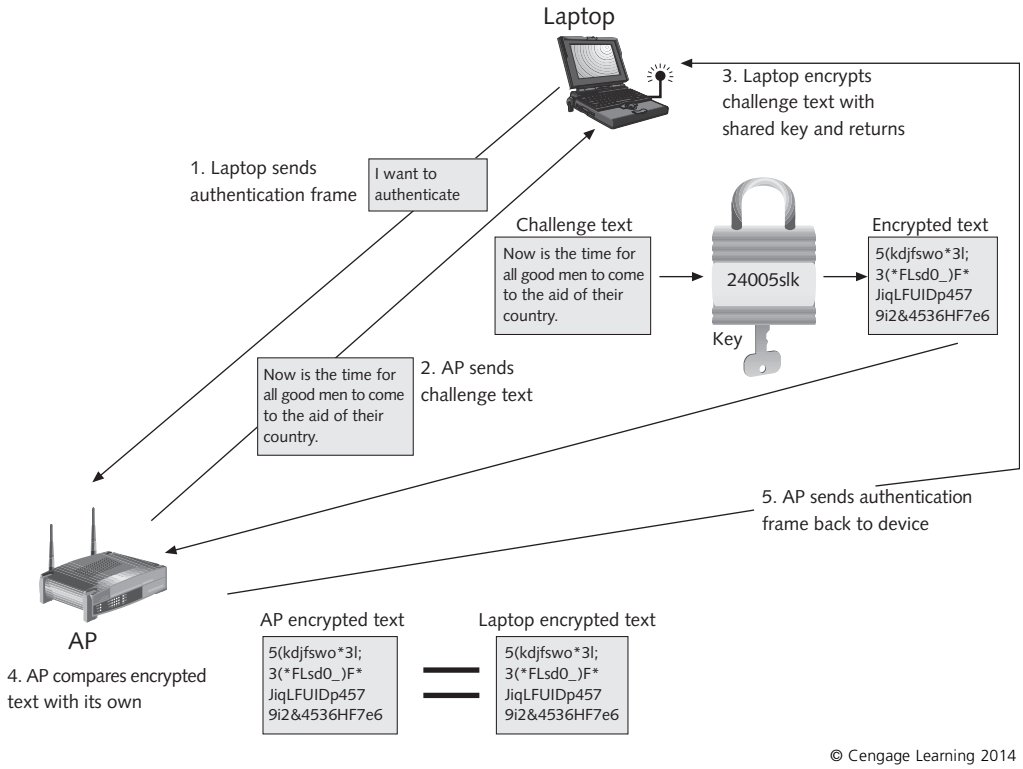

© Cengage Learning 2014

Figure 7-6 Shared key authentication

Shared key encryption in WEP uses the Rivest Cipher version 4 (RC4) encryption algorithm, which allows keys of up to 128 bits. The 802.11 standard uses a 40-bit or 104-bit key with a 24-bit **initialization vector (IV)** added to the beginning of the key. The IV initializes the key stream generated by the RC4 algorithm and is transmitted in cleartext.

To understand WEP's weaknesses, you need to understand IV vulnerabilities. The IV is part of the RC4 encryption key, and because it is transmitted in cleartext, it gives attackers 24 bits of the key. The IV is a short stream by cryptographic standards, and reusing the same key results in keys that repeat after a short time. All users use the same key, so capturing enough packets to crack the key and decrypt transmissions is not difficult.

Often, the IV starts at 0 and increments by 1 each time a key is generated, so not only does it repeat, it does so in a predictable pattern. Because the IV is a 24-bit value, only 16 million combinations are possible. Therefore, in a busy network, the IV can reinitialize and start over at 0 in about 6 hours. After enough packets have been captured, attackers can crack the key with a brute-force or dictionary attack.

Because the challenge text in shared key authentication is sent in cleartext, attackers can capture it along with the IV and then capture the challenge response. With this information, they can crack the key. Even though the RC4 specifications make it clear that keys should never be reused, inevitably they are in WEP.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

WEP provides adequate protection against casual users with no ill intent, but not against attackers who are determined to gain access. Dynamic WEP, a newer version, offers slightly better protection because it rotates keys frequently, which solves the IV problem. It also uses different keys for broadcast and unicast traffic and changes them frequently. Dynamic WEP requires minimal effort to set up, but it is not used much because it provides only a partial solution. Another option is WEP2, which was developed to address some vulnerabilities in WEP. WEP2 uses a 128-bit key and Kerberos authentication. Although these enhancements help, WEP2 is no more secure than WEP and is not used widely.

WEP is often disabled to increase throughput. Even though it does not provide much security, you should still make sure it is enabled as a first line of defense to prevent casual users from connecting.

Default WEP Keys

Even though the 802.11 standard states that APs and wireless stations can hold up to four keys simultaneously (see Figure 7-7), only one is chosen as a station's **default key** to encrypt messages for transmission. The default key does not have to be the same on every station, but the same key must be used for both encryption and decryption. Each station must contain the key used for encryption to decrypt messages. For example, if a sending station uses Key 3 for encryption, the receiving station must use its copy of Key 3 to decrypt the message.

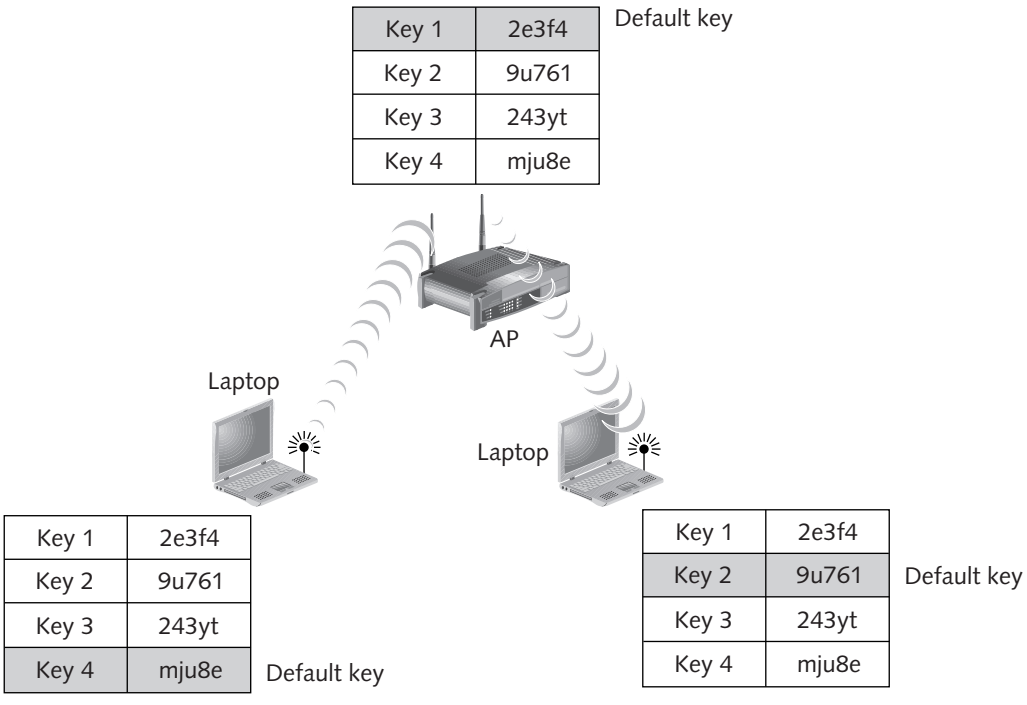

Figure 7-7 Default WEP keys

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Key Management Concerns in 802.11 Networks

Key management is a challenge in wireless security. The 802.11 standard leaves the details of key management up to vendors and users.

WEP was integrated into devices to prevent casual eavesdropping. It works well to prevent eavesdropping but has shortcomings in preventing unauthorized access. Furthermore, WEP keys must be installed on all stations in a network before they can communicate. This procedure takes a lot of time, which often means that keys are changed infrequently or not at all. WEP and WEP2 should not be used for security.

Even if stronger encryption methods are used, an effective key management method is still crucial. In wired networks, many organizations use some form of Public-key Infrastructure (PKI). If a comprehensive PKI system is already in place, it should also be used for wireless segments to make sure that keys are managed effectively.

MAC Address Filtering and Spoofing

Wireless networks depend on MAC addresses for more than establishing, maintaining, and managing connections or transmitting data. Wireless stations also use MAC addresses rather than IP addresses for identification between stations and APs. IP addresses are Network layer identities and cannot be used for addressing in a MAC frame.

MAC addresses are hard-coded into NIC firmware, but entering a different value into the frame is a simple matter with the right software tool. You can use most configuration tools that are included with hardware to change a WNIC's MAC address, and you can find many free tools that are specifically designed for spoofing addresses.

One of the most basic security mechanisms in wireless networks is MAC address filtering. MAC addresses of legitimate stations are entered manually in the configuration interface of the AP's MAC address table, and only recognized stations can connect to the AP. However, attackers get around this mechanism by sniffing the network to find legitimate MAC addresses, and then they alter frames by using those values. This tactic is called MAC address spoofing. The injected frame appears to be from a legitimate source, so attackers can connect to the AP and network.

Wireless Device Portability

Mobility is the primary purpose of wireless devices, such as laptops, smartphones, and tablets, but this capability makes them vulnerable in many ways. Wireless devices are designed to be portable, which makes them vulnerable to theft, unauthorized use, improper or unsafe storage and handling, established connection protocols being bypassed, and more.

Ensuring the physical security of a company-owned wired computer is relatively simple; it stays on a user's desk unless authorized personnel move it. If someone else moves or tampers with the computer, measures are in place to spot these events quickly, locate the device, and deal with the problem. Other security measures include preventing unauthorized entry to the facility, protecting established communication channels between the corporate network and Internet, scheduling and installing software updates automatically, and ensuring that devices remain within the corporate security perimeter, under the IT Department's control.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Mobile devices are a different story. Employees take these devices outside the office and use them to connect from home, their cars, and hotel rooms, and they might allow others to use the devices. Theft is a problem, too. For example, in 2006 an employee of the U.S. Department of Veterans Affairs had a laptop stolen from his car. The laptop contained the personal information of more than 50 million veterans, all of whom had to be notified, even though the laptop was recovered and showed no signs of tampering. However, these veterans still had to worry and watch for signs of identity theft because of the possible breach of a single mobile device.

Also, mobile devices might not be backed up properly or have updates installed if they are not connected to the corporate network when those procedures take place. Mobile devices can be compromised in countless ways when they are not protected by security personnel. For this reason, make sure that highly sensitive data is not stored on mobile devices, and ensure that the devices use strong encryption and authentication.

Examining Wireless Security Solutions and Countermeasures

In the early years of wired networking, wireless standards focused not on security but on connectivity. Therefore, wireless security has lagged a few years behind wired network security, but the IEEE, Wi-Fi, and other standards organizations have worked to address the inherent security problems in wireless networks.

After learning about 802.11 security issues, you might think that the best option is simply to avoid using wireless networks. For some businesses, the risks of wireless are too great to justify its use, and many organizations do not want the headaches of securing wireless networks. However, wireless networking offers too many advantages to abandon it just because security is challenging. In the following sections, you learn some common solutions for addressing the security flaws in wireless networks. You also learn about special security requirements of wireless networks and common configurations that mitigate wireless vulnerabilities and protect against wireless networking threats. With a little extra work, wireless can be a secure, effective enhancement to networking technology.

Incorporating a Wireless Security Policy

A key component of security is a wireless security policy. The main differences between wireless and wired networks are in the Physical and Data Link layers of the OSI model, so your wireless security policy should focus on them. A wireless security policy should address the following issues:

- Scope and goals of the policy
- Responsibilities for wireless matters and contact information for responsible parties
- Physical security of APs to prevent tampering, unauthorized access, interference sources, and other problems

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Approved hardware and software
- Procedures for requesting, testing, installing, and configuring hardware and software
- Assignment of responsibilities for installing, maintaining, and managing wireless devices
- Guidelines and penalties for scanning or accessing the wireless network without authorization
- Explicit statements about the nature of wireless communications, including measures to protect the rest of the network from potential harm from wireless communication
- Details on wireless security awareness training
- Internet access via wireless connections
- Assignment of responsibilities for protecting data, privacy, and devices
- Penalties for attempting to bypass security measures willfully

The policy should also address other details related to the wireless environment, such as requirements for encryption methods, authentication, and storage of confidential data.

Ensuring Physical Security

Just as no computer can be made completely secure, you cannot ensure absolute physical protection for mobile devices. The best tool for ensuring physical security is to provide security awareness training for users. They should be made aware of the potential for theft and the consequences of stolen devices, be trained not to leave their mobile devices where they can be stolen or used by unauthorized people, and be trained not to leave wireless devices logged on to the network.

Training should also include instructions for protecting mobile devices from damage, such as making sure that mobile devices are stored in an appropriate environment. Remind users that some protective measures are simply common sense; for example, they should not leave laptops in cars during the summer or winter, and should never leave laptops unattended in public under any circumstances.

Because of the increasing frequency of data loss and property loss associated with employee misuse of mobile devices, some organizations have revised their policies to make employees liable for such losses. In addition, some organizations limit the type of company data that can be accessed from mobile devices or stored on them.

Planning AP Placement

You can avoid many pitfalls of wireless security with careful planning. Because RF signals propagate in different directions at different rates and distances, and signals are affected by interference sources, determining the exact signal direction is difficult. A **site survey** is a procedure for assessing the environment and determining where APs are needed to provide adequate coverage. Site surveys also help determine whether to use directional or omnidirectional antennas and how much power is necessary for required signal performance.

Most important from a security standpoint, a site survey tells you if your signal extends beyond areas that are within your physical control. Network components require careful

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

placement to provide adequate coverage yet prevent indiscriminant radiation of the signal. Such placement can mitigate the risk of many vulnerabilities; attackers cannot detect and infiltrate your network if they cannot detect a signal.

A site survey follows these general steps:

1. Get facility blueprints and other diagrams that show walls, walkways, wiring, elevators, other WLANs if applicable, and planned locations of wireless stations.
2. Examine the facility for potential sources of interference with the RF signal, such as metal racks, power conduits, and microwave ovens. Be sure to check the diagram for accuracy and note any discrepancies.
3. Determine the coverage areas needed, identify roaming areas, and specify where there will be no coverage. You can limit potential problems by limiting roaming areas.
4. Determine preliminary AP locations. There should be some overlap in coverage, but make sure that the channels are far enough apart to avoid interference.
5. Identify suitable mounting locations for APs, antennas, patch cabling, and power lines, and consider different antenna needs based on location. For example, you do not want to use an omnidirectional antenna for an AP mounted on or near an outside wall. A more suitable choice would be a high-gain directional antenna oriented to transmit and receive within the facility.
6. Test AP placement before running the wiring and mounting APs permanently. Several RF site survey tools, such as AirMagnet, are available that can identify the AP, its signal strength and quality, and its data rate. Remember to test coverage on lower and upper floors of a multifloor facility.
7. Walk the entire perimeter outside the facility to discover whether any signals are leaving the building. If you find any, identify exactly where they go. Check upper and lower floors in the building for signals as well. You might be unable to prevent all signals from radiating outside, but you need to assess the potential risk. For example, you might have an RF signal that radiates 10 feet outside the office, but it might be in an area that is not easily accessible to attackers or is secured by external fencing or security patrols.
8. Document the results of the site survey. On the facilities diagram, note each AP's location, signal coverage, and antenna type. This information also helps you identify rogue devices and aids in troubleshooting.

Changing Default Hardware and Software Settings

Default settings for hardware and software in wireless configurations create serious vulnerabilities, so you should change a number of settings before installing a new device. For example, most APs include a simple HTTP interface for configuration, but after you configure the AP, you should disable the interface to prevent attackers from being able to access the AP and tamper with it. If you do not want to disable HTTP access, you can add password-protected access to the interface. However, disabling the HTTP interface is the preferred method.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

You should change the following default settings before connecting to the network:

- Change the default SSID so that it does not give attackers information about the type of equipment used or its location. Default SSIDs commonly include information about a device's manufacturer, as you saw with the "linksys" SSID used for Betty's wireless router in the rogue device example earlier in this chapter. Attackers who discover a default SSID can assume that other default values are being used, such as the administrator password. Attackers can then use this information to attack the network.
- Change the default administrator password to be as strong as the AP can support.
- Change the default beaconing interval, primarily to reduce traffic and improve throughput for data.
- Manufacturers' default keys are well known, too. When configuring a device, specify another key.
- Retaining default channels can increase the likelihood of co-channel interference, which might reduce performance or disable the network entirely. Also, attackers frequently scan for vulnerable wireless networks by using known default channels. If default channels are not changed, attackers might discover the device vendor's information and other default values they can use in attacks.
- By default, most security measures are not enabled on wireless equipment. Before placing a device on the network, configure it with appropriate security settings, including MAC ACLs, authentication, and encryption.

As you can see, changing default values makes it harder for attackers to find a network and gain access to it.

Strong Encryption and Authentication

Strong authentication and encryption were not addressed in the original IEEE 802.11 standard, but newer methods have been developed to improve security. These methods are discussed in the following sections.

802.1x and Extensible Authentication Protocol

The **IEEE 802.1x** standard was developed to provide port-based access control on Ethernet LANs, and it was revised to work for wireless networks. The standard uses **Extensible Authentication Protocol (EAP)** and encrypted tunnels for data exchange. EAP is a group of extensible management protocols that stations use to request port access, and it includes a method of secure key exchange.

802.1x authentication involves three main participants: the supplicant, authenticator, and authentication server (see Figure 7-8). The supplicant is the station that requests access through the authenticator, which is usually an AP in wireless networks. The authenticator passes the request to the authentication server, which stores the credentials of authorized users. An important strength of 802.1x is that the supplicant never communicates directly with the authentication server, thus reducing the chance of compromising sensitive data stored on the server.

Often, a RADIUS server is used for the authentication server. When the authentication server receives a request, it checks to see whether the authenticator is authorized to send

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

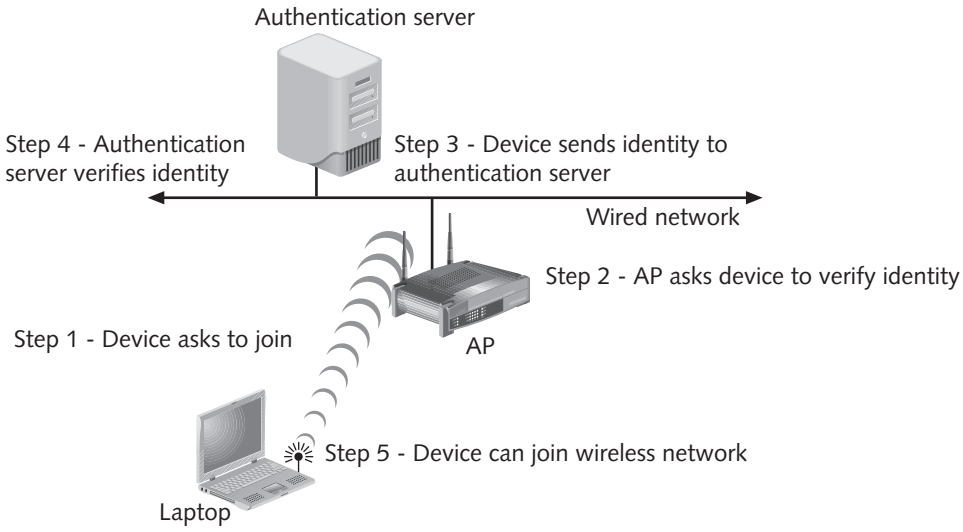

© Cengage Learning 2014

Figure 7-8 802.1x authentication

requests. If so, the server checks its database for the user's name. If it finds the user's name, it checks the password to determine whether to grant access. Depending on the authentication method, the authentication server might respond with a challenge message to which the user must respond correctly. The authentication server notifies the authenticator of its decision, and the authenticator passes the decision back to the supplicant. The authenticator relays all messages between the authentication server and supplicant.

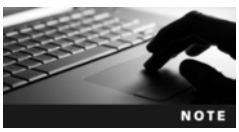

Although not defined by standards, wireless user credentials can also be stored in an external database, such as SQL, LDAP, or Active Directory.

The 802.1x standard is based on EAP, which includes the following extensible protocols:

- *EAP-Transport Layer Security (EAP-TLS)*—Requires digital certificates to validate supplicants.
- *EAP-Tunneled TLS (EAP-TTLS)*—Supports advanced authentication, such as tokens. It does not require client-side digital certificates, which makes it more popular than EAP-TLS, but slightly less secure.
- *Protected EAP (PEAP)*—Uses certificates from supplicants but does not require certificates from the authentication server. After authentication, the authenticator builds an encrypted tunnel to authenticate the supplicant.
- *Lightweight EAP (LEAP)*—This Cisco proprietary standard authenticates based on the Windows username and password. No certificates are required. This standard is not considered secure, even by Cisco, which recommends using newer protocols such as FAST or PEAP.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *Flexible Authentication via Secure Tunneling (FAST)*—Considered by Cisco as a replacement for LEAP, FAST sets up a tunnel without checking digital certificates and supports tokens.

802.11i and Advanced Encryption Standard Ratified in 2004, the **IEEE 802.11i** standard addresses encryption and authentication. It uses 802.1x authentication and **Advanced Encryption Standard (AES)**, a symmetric block cipher, for encryption. AES is strong enough to meet the U.S. Federal Information Processing Standard (FIPS).

Unlike a stream cipher that substitutes character for character, a block cipher breaks data into blocks of 8 to 16 bits, and then encrypts each block separately. For additional security, blocks can be arranged randomly rather than sequentially. Several iterations (rounds) of complex multiplication operations are performed on each block to substitute and rearrange bytes, and then to encrypt the new arrangements of bytes again. The number of rounds depends on the key length. With a key length of 80 bits, breaking the AES key with a brute-force attack would take 800 years and $1 million worth of computers. At 256 bits, breaking the AES key would take 10^{56} years. That's 10 followed by 56 zeros, which is a *very* long time, but attackers will probably discover how to break the AES key sooner than that. Although AES is considered secure for now, it is reasonable to expect that security experts and attackers will discover exploitable flaws as technical capabilities continue to advance.

Wi-Fi Protected Access **Wi-Fi Protected Access (WPA)** was developed and released by the Wi-Fi Alliance during the same period that the IEEE was working on the 802.11i standard. WPA replaced WEP encryption with **Temporal Key Integrity Protocol (TKIP)**. TKIP is based on the same RC4 mechanisms used in WEP, but it includes a method for generating new keys for each packet, which addresses WEP's problem of static keys that repeat. TKIP also incorporates automatic AP-client synchronization of unicast encryption keys on a frame-by-frame basis, so keys do not have to be changed manually, as they do in WEP.

In TKIP, different keys are used for authentication and encryption. Keys used between a pair of stations are called **pairwise keys**. After a station's credentials are accepted, the authentication server produces a master key for that session, and TKIP distributes this key to the station and AP. This **pairwise master key (PMK)** generates data encryption keys, data integrity keys, and session group keys for multicasts. The PMK is not used for any other purpose, which makes it less vulnerable to attack.

At the start of a session, the PMK is used to compute keys for encryption, data integrity, and authentication. The **pairwise transient key (PTK)** is the first key created from the PMK. It is actually four keys that are shared between the AP and wireless client. PTKs are used to generate more keys according to the cryptographic algorithm in use, whether it is TKIP or AES. Keys are transient, meaning they change in real time during a session, and a new key is generated dynamically for each packet.

WPA also uses **Message Integrity Check (MIC)**, a mathematical function similar to the cyclic redundancy check (CRC), to check messages for evidence of alteration (see Figure 7-9). The MIC implementation in TKIP is referred to as Michael, and it uses a longer IV of 48 bits. If the MIC value for a received message does not match the value sent with the message, corruption is assumed and the packet is dropped. Unlike CRC, however, MIC includes an

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

optional measure that deauthenticates a station and prevents new associations for one minute if MIC errors are detected. MIC also includes a frame counter that prevents replay attacks.

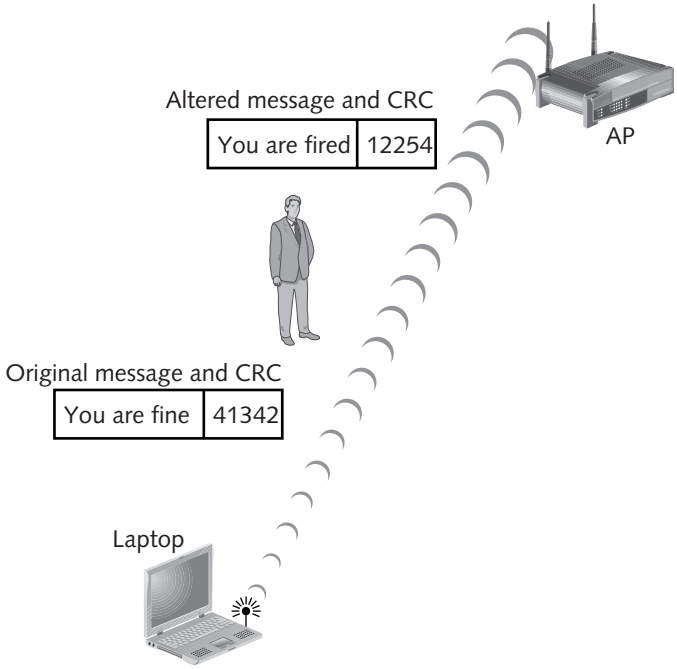

© Cengage Learning 2014

Figure 7-9 The MIC process

WPA authentication can work with 802.1x or a preshared key (PSK). If you use a preshared key, it must be entered on the AP and station first and used as a seed for generating encryption keys. A major disadvantage of using a passphrase and preshared keys is that the passphrase must be more than 20 characters and very complex—in other words, it cannot use dictionary words and it must consist of letters, numbers, and symbols. Otherwise, the passphrase is susceptible to being cracked.

WPA offers the following notable improvements over WEP:

- Minimum key length is increased for stronger data encryption.
- IV sequencing is enforced so that IVs are not reused.
- IV length is doubled from 24 bits to 48 bits.
- Packet-tampering detection is built into MIC.
- Key rotation is automatic.

Wi-Fi Protected Access version 2 (WPA2) is based on the final ratified 802.11i standard; WPA2 uses AES for encryption and uses 802.1x or preshared keys for authentication. One notable difference between 802.11i and WPA2 is that WPA2 allows both TKIP and AES clients to communicate, whereas 802.11i recognizes only AES. In 2010, the Wi-Fi Alliance

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

extended the certified EAP types beyond EAP-TLS, which had been the only certified type to that point.

You can use WPA and WPA2 in Personal Security and Enterprise Security modes of operation, which provide guidance on using WPA and WPA2 for different purposes. Personal Security mode is designed for single users or small office/home office (SOHO) settings, and Enterprise Security mode is designed for medium-sized to large businesses.

Table 7-3 summarizes the major wireless security solutions.

Type	Encryption	Authentication	Key length	Security level
WEP	WEP	Shared key	40-104 bits with a 24-bit IV	Low
WEP2	WEP2	Kerberos	128 bits	Low
Dynamic WEP	Dynamic keys	Various	128 bits	Low
WPA Personal Security	TKIP	PSK	128 bits	Medium
WPA2 (802.11i) Personal Security	AES	PSK	256 bits	Medium to high
WPA Enterprise Security	TKIP	802.1x	128 bits	Medium
WPA2 Enterprise Security	AES	802.1x	256 bits	High

© Cengage Learning 2014

Table 7-3 Wireless security solutions

Recent research has demonstrated serious weaknesses in WPA and WPA2 when using TKIP. In 2008, the Beck-Tews attack demonstrated that attackers could inject 7 to 15 packets of their choice into a network. This attack involves the use of waiting periods to circumvent the 60-second MIC limitation noted in the preceding WPA section and the use of Quality of Service (QoS) channels. In 2009, Finn Michael Halvorsen extended this attack to allow the injection of larger packets. Denial of service and ARP poisoning are among the attacks that are possible based on the Beck-Tews exploitation technique. In 2010, Beck described another attack that allows all packets to be decrypted as they are sent to a wireless station. Because of these recent developments, WPA2-TKIP is now considered far less secure than WPA2-AES.

WPA2-AES Enterprise Security provides the highest security available and should be used to protect a wireless network when an authentication server is available. (An authentication server is required for 802.1x authentication to work.) If an authentication server is not available, use WPA2-AES Personal Security instead.

In 2007, the Wi-Fi Alliance released **Wi-Fi Protected Setup (WPS)**, a protocol designed to automate key distribution in small office and home networks. Instead of requiring users to enter long passphrases in each device, one option allowed users to enter an eight-digit PIN.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

In late 2011, a design flaw was discovered that allowed brute-force attacks against the PIN and subsequent access to the network it was supposed to protect. At this time, the only defense against this flaw is to disable WPS if possible.

Wireless Auditing

Auditing wireless networks is an integral part of security management, and it can be more important than auditing wired networks. Ensuring that only authorized users can access data is critical in wireless networks, which are more vulnerable to interception. Audits are based on security policies, so auditors review the security policy first and then test the system to ensure compliance. Next, they assume the role of attackers to check for vulnerabilities. You can search the Internet and easily find all the tools and information needed to attack any network; many of these resources are free.

Security auditors are often called "penetration testers" or "ethical hackers."

Hiring third-party experts to conduct security audits can be a good idea for the following reasons:

- They see your network with fresh eyes and no preconceived ideas. When you have worked with your system for a long time, you might miss even obvious problems.
- They are likely to have different skills and an extra arsenal of tools and techniques.
- They usually do more wireless security audits than administrators, so they have the focus and experience of a specialist.

If you hire a third-party expert, make sure to check credentials and ask for references; a few unscrupulous people pretend to be qualified, trustworthy experts. Also, make sure that responsibilities, fees, and permissions are spelled out in a contract. Genuine auditors almost always require one.

Risk and Security Assessments An early step in developing a security policy is to create a comprehensive risk assessment. You must know what your assets are and how critical they are to your organization before you can determine how to protect them appropriately. Risk assessment helps with this critical task.

Most basic security audits do not include a comprehensive risk assessment; audits are intended to evaluate policy compliance, determine the effectiveness of existing security, and identify unknown vulnerabilities. A risk assessment starts with an inventory of company assets, and is followed by an analysis of possible threats to those assets, the consequences to the organization if a threat materializes, the probability that the threat could occur, the security controls available to mitigate the risk, and the organization's acceptable level of risk. A security assessment differs from a risk assessment in that it identifies existing security measures. To upgrade security and possibly improve performance and function, you need to evaluate the assessment's starting point accurately. The security assessment is then used during a security audit to evaluate the effectiveness of your security measures.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Auditing Tools After an auditor completes a security assessment, the next task is to use a variety of tools and techniques to attempt to break into the network. This task, formally known as **penetration testing**, is intended to identify security vulnerabilities that attackers could exploit. For example, attackers use sniffers in the network reconnaissance phase to capture packets so that they can gather information about targets. Auditors use sniffers to see what kind of information attackers can gain by using them. Table 7-4 describes several popular wireless sniffers and the types of information they can gather. Hundreds of sniffing programs are available for PCs, handheld devices, and any available OS, and many add-ons can be combined with these programs.

Name	OS	Description
Aircrack-ng	Mac OS X, Windows, Linux	A sniffer and WEP/WPA key cracker that can be used to launch attacks; consists of airodump-ng (packet capture), airdecap-ng (decryption of WEP/WPA capture files), aircrack-ng (WEP and WPA/PSK cracker), and aireplay-ng (packet injection)
Kismet	POSIX, Linux, BSD, Mac OS X, Win32	An 802.11 Layer 2 detector, sniffer, and IDS that is passive and undetectable in operation; can detect hidden networks, rogue devices, attacks, and unauthorized users
NetStumbler	Windows	An 802.11a, 802.11b, and 802.11g WLAN detector and analyzer; WebStumbler is an add-on that converts files to HTML
Wireshark (previously Ethereal)	Windows, UNIX, Linux, and other OSs and UNIX variations	A network protocol analyzer capable of live monitoring or capture; live data can be read from FDDI, PPP, token ring, Ethernet, and 802.11 networks
AirSnort	Windows, Linux	A tool that recovers encryption keys by using passive monitoring; requires 5 to 10 million packets
FakeAP	Linux, BSD UNIX	A tool that generates thousands of counterfeit 802.11b APs and is useful as a honeypot tool
WPA Cracker	Linux	A tool that prevents dictionary and brute-force attacks against WPA

© Cengage Learning 2014

Table 7-4 Wireless sniffers

AP Logging Functions

A network that cannot be monitored cannot be secured, so the ability to maintain and review a record of activity is essential for security. Less sophisticated APs might not be able to generate more than simple connection statistics and a basic event log, but many enterprise-class models can maintain more complex event logs and connection statistics, and can interface with a Simple Network Management Protocol (SNMP) tool. Figure 7-10 shows an AP event log.

SNMP requires the installation of an SNMP agent on the device you want to monitor (in this case, the AP). Logged information is stored in the SNMP agent's management information base (MIB), and it is housed on the monitored device as well. The MIBs gather data and

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

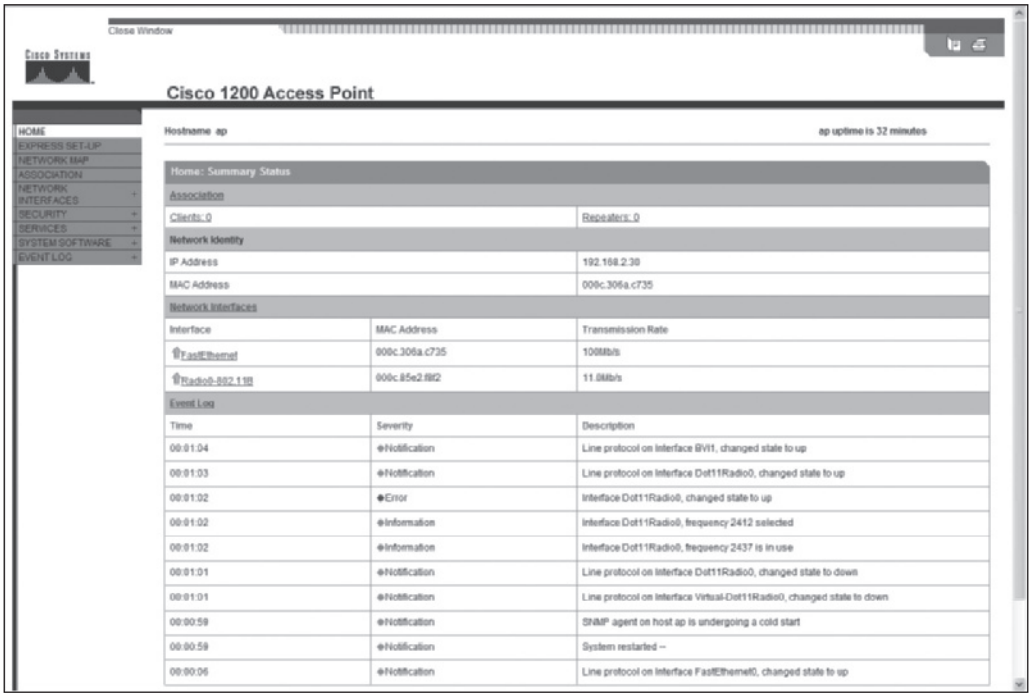

Source: Cisco Systems

Figure 7-10 An AP event log

report it to an SNMP management station that produces overall network statistics. An SNMP configuration is shown in Figure 7-11. Using the network statistics, you can set an SNMP alarm that sends an alert message, called an SNMP trap, to the management station when specified limits are exceeded. The management station then queries all stations for details about the event that triggered the alert.

Best Practices for Wireless Network Security

You have a number of options for protecting wireless devices. This section summarizes some best practices, but the list is not comprehensive. First and most important, make sure that all wireless device users have received thorough training on corporate security policies. Second, restrict the data that can be transferred to a portable device. Highly sensitive data should never be stored on a portable device unless it is absolutely necessary. Third, try to break into your wireless network. You can use the same tactics and tools that attackers use, and your knowledge of the network's configuration might allow you to probe for vulnerabilities that attackers do not know. If your budget allows it, a third-party wireless security audit is a sound investment.

You must have written permission from an authorized party before performing penetration testing on any network. Attorneys should review all contract terms, including work specifications, possible effects, timelines, and expectations of both parties.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

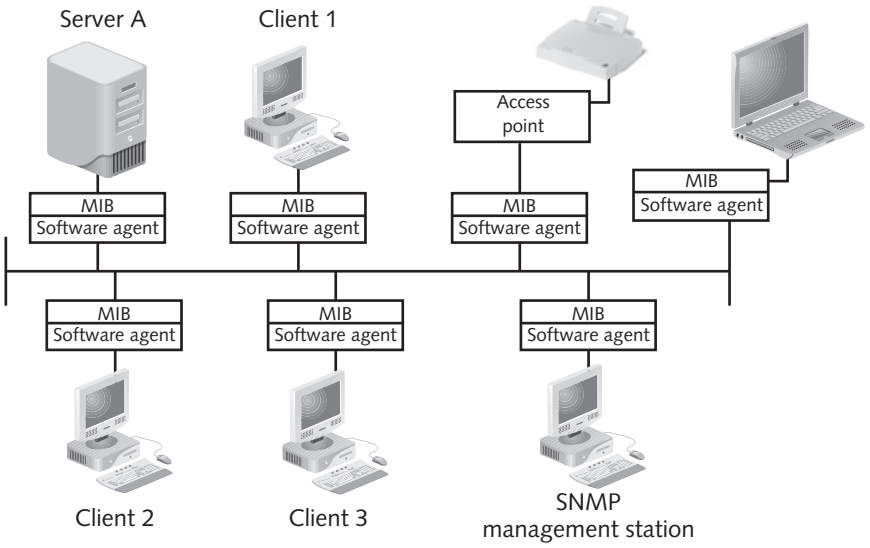

© Cengage Learning 2014

Figure 7-11 SNMP configuration

The following list summarizes other best practices to follow:

- Use strong authentication, such as 802.1x.
- Use strong encryption, preferably end to end.
- Perform a site survey and place APs strategically.
- Make sure that a comprehensive wireless security policy is kept up to date, and that users are trained to follow it.
- Change default settings, such as SSIDs.
- Avoid using protocols that send traffic in cleartext.
- If appropriate, use VPNs for wireless transmissions.
- Use wireless IDPSs.
- Make sure that all stations use updated antivirus protection.
- Make sure that wireless devices use firewalls.
- Audit the wireless network periodically.
- Monitor your wireless network traffic with the best tools available.

Mobile Device Security

Many modern IT departments resisted the addition of mobile wireless devices to their areas of responsibility. While the integration of laptops was unavoidable, many chief information officers and chief security officers were reluctant to include handheld devices under their supervision. Cellular phones were thought to be the purview of telephone service providers, and the security challenges associated with cellular phones and PDAs (personal digital assistants) were intimidating. However, as cellular phones became smartphones that could access

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

the Internet and use mobile applications for business activities, the pressure from company management increased and mobile devices finally had to be added to the corporate network.

Thus, IT security analysts faced the problem of securing handheld devices and the transmissions to and from the devices. Many of the same threats, risks, and consequences associated with corporate network security apply to the security of corporate mobile devices. However, there are important differences. Traditional corporate assets such as servers, workstations, and transmission media are under the company's physical control, but handheld wireless devices are often outside the physical control of the IT security team. Also, the transmission media used by these devices might be beyond a company's control; these media include the company's WLAN, the network of the telephone service provider, the Internet, hot spots, and private WLANs. In addition, users may synchronize their handheld devices with computers that are not controlled by the corporate IT department, which increases the risk of malware infection. These problems can then spread from the handheld device to the corporate network.

Today's popular handheld devices include smartphones and pads. A pad is a tablet-like device that is substantially larger than a cell phone but smaller than a laptop. These handheld devices typically use three types of network connections: cellular, which is controlled by the telephone service provider; wireless, which is based on IEEE 802.11 and controlled by the party that controls the WLAN; and Bluetooth, a short-range radio frequency standard used for wireless earpieces. This flexibility provides great benefit to users but poses numerous security problems for the corporate security team.

Loss of corporate data is also a major concern. Not only could proprietary information fall into the hands of competitors, data loss could make a company liable for violations of statutes such as FERPA (Family Educational Rights and Privacy Act), HIPAA (Health Insurance Portability and Accountability Act), and SOX (Sarbanes-Oxley Act). In some extreme cases, national security could be affected.

This chapter has already addressed some of the threats to mobile devices; for example, a user might lose a device or have it stolen. Additional threats include the unintentional installation of malware via a Trojan horse attack, packet sniffing and electronic eavesdropping through spyware that records voice and data activity, unauthorized access of corporate networks, and exploitation of GPS (Global Positioning System) capabilities to track a user's location.

For several reasons, handheld devices have suffered fewer attacks than traditional workstation systems. For example, handheld devices do not use a predominant operating system (OS), although several companies have joined the competition, including Android, Symbian, Apple iOS, Research In Motion (BlackBerry), Bada, and Microsoft. In addition, cellular network providers have tended to maintain a closed-systems approach in their networks, applications, and devices. However, attacks against handheld devices will probably increase as the battle for OS market share narrows down the software targets and more open standards are implemented to facilitate interoperability among devices.

Approaches to Mobile Device Security

As with all network management, you should carefully consider the topic of mobile device security before implementation rather than after the fact. When determining how to implement security controls for mobile devices, your decisions should consider asset inventory,

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

risk assessment, and asset prioritization based on business process analysis. In addition, these decisions should directly reflect the standards set by the organization's security policy. A checklist that ensures the security of handheld devices should include the following items:

- Device configuration management
- Critical patch and OS update management
- Application installation and configuration management
- Elimination of unneeded applications and services, such as Bluetooth and Wi-Fi
- Antivirus software
- Firewall software
- IDPS software
- Antispam software
- Antispyware software
- Remote content erasure capability
- Remote password reset capability
- VPN software
- Backup management
- Authentication management
- Encryption
- Log management
- Incident response policy and procedures
- Restriction of application downloads
- Restriction of camera, microphone, and removable media use
- Remote diagnostics
- Subscriber Identity Module (SIM) security
- User training

Many of these controls are also typical in wired networks, but implementing the controls in a wireless environment presents special challenges. A number of products are available to help you control and manage security problems in an enterprise environment that includes handheld devices.

Chapter Summary

- A major challenge for wireless networking is security. Vendors, developers, and the IEEE have worked to come up with solutions to wireless security threats.
- Wireless networks use the airwaves as a transmission medium, so packets are more vulnerable to interception.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- The Media Access Control (MAC) sublayer of the OSI model performs many critical functions in a wireless network. Each access point (AP) has an SSID for a network name and communicates by using management, control, and data MAC frames.
- Passive scanning involves the use of a wireless station that listens for beacon frames. A passive attack uses passive scanning to gather information for later use. Active scanning involves the use of a station that sends probe request frames on each available channel and waits for a response.
- Before a station can be associated with an AP and join a wireless network, it must be authenticated. The IEEE 802.11 standard includes open system authentication and shared key authentication.
- SSIDs and other information about the wireless network are vulnerable in standard 802.11 transmissions because management frames send network information in cleartext.
- WEP was implemented in the original 802.11 standard and uses the Rivest Cipher version 4 (RC4) encryption algorithm. A default key is used for encryption. The initialization vector in a WEP key is vulnerable to being cracked.
- Even though wireless networking has inherent security problems, a number of effective security solutions and standards are available, including IEEE 802.1x, WPA/WPA2, and IEEE 802.11i. Recently, experts have found that WPA/WPA2 using TKIP is vulnerable to attack.
- Auditing a wireless network is crucial to maintaining and improving security. Security assessments, third-party audits, penetration testing, and logging are part of a comprehensive auditing plan.
- Less sophisticated APs might generate only simple logs, but many enterprise-class models can maintain an event log and connection statistics, and can interface with a Simple Network Management Protocol (SNMP) tool.
- Some best practices for wireless security include training users, developing a wireless security policy, restricting the data stored on portable devices, and ensuring that default settings are changed.

Key Terms

acknowledgement (ACK) A type of control frame that a receiving station sends when a packet is received successfully with no errors.

active attacks Attacks that attempt to gather information for subsequent attacks by using active scanning—sending probe request frames on each available channel.

active scanning The process of a wireless station sending a probe request frame on each available channel and waiting for a probe response frame from available APs.

Advanced Encryption Standard (AES) A symmetric block cipher used for encryption in IEEE 802.11i.

association A two-step process of being accepted into a wireless network. First, a station listens for beacon frames to locate a network to join, and then it goes through the

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

authentication process. Second, the station sends an association request frame; if it is accepted, the station receives an association ID and connection information. A station can be associated with only one network at a time, and it must be authenticated before being associated.

clear to send (CTS) A type of control frame sent in response to a request to send (RTS) frame; it gives the sender clearance to begin transmitting packets.

control frame A type of MAC frame used to help deliver data frames between stations and to control access to the medium.

data frame A type of MAC frame that contains the TCP/IP packet sent over a wireless network or between wireless devices. It carries data from higher-layer applications, such as printer control data or Web pages.

default key A wireless station or AP's key for encrypting messages for transmission. The default key does not have to be the same on every station, but the same key must be used for both encryption and decryption.

Extensible Authentication Protocol (EAP) A group of extensible management protocols used in IEEE 802.1x. EAP includes a method of secure key exchange, and wireless stations use it to request port access.

IEEE 802.1x A standard developed to provide port-based access control on Ethernet LANs; it has been revised for wireless networks to incorporate authentication and key management.

IEEE 802.11 authentication The process by which an AP accepts or rejects a wireless device's connection attempt.

IEEE 802.11i A wireless security standard intended to replace the IEEE 802.11 WEP-based standard. It uses 802.1x authentication and AES for encryption. See also *Advanced Encryption Standard (AES)*.

initialization vector (IV) In WEP's shared key encryption, a 24-bit value added to the beginning of a key to initialize the key stream generated by the RC4 algorithm.

management frame A type of MAC frame used to establish and maintain communications between wireless devices, or between a wireless device and an access point.

Message Integrity Check (MIC) A mathematical function used in WPA that replaces the cyclic redundancy check (CRC); MIC is designed to detect tampering in packets. See also *Wi-Fi Protected Access (WPA)*.

open system authentication An 802.11 authentication method that relies on a station having an SSID that matches the network it is attempting to join; a matching SSID is the only criterion for acceptance. The security drawback of open system authentication is that SSIDs are transmitted in cleartext in management frames.

pairwise keys Any keys used between a pair of devices in TKIP. See also *Temporal Key Integrity Protocol (TKIP)*.

pairwise master key (PMK) The TKIP key used to generate data encryption keys, data integrity keys, and session group keys, among others. This key is used only once at the start of a session.

pairwise transient key (PTK) The first TKIP key generated by the pairwise master key (PMK) and used for further key generation, according to the cryptographic algorithm in use.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

passive attack An attack that uses passive scanning to gather information for later use in other attacks; also called network reconnaissance.

passive scanning The process of a wireless station listening on each available channel for an AP's beacon frame, and then moving on to the next channel without sending anything.

penetration testing The process of using a variety of tools and techniques to attempt to break into a network. Penetration testing, also called ethical hacking, is used legitimately as part of security audits to identify security vulnerabilities that attackers could exploit.

power-save poll (PS-Poll) A type of control frame that a station sends after awakening from power-save mode; the frame indicates to the access point that it is ready to receive any frames the AP has buffered for it.

radio frequency (RF) monitor mode A wireless NIC mode, equivalent to promiscuous mode in wired NICs, that allows a WNIC to capture packets without authenticating or associating with an AP or ad-hoc wireless network.

request to send (RTS) A type of control frame that a station sends when it wants to transmit. The RTS frame sender requests that the medium be reserved long enough for the transmission to be completed and for an acknowledgement control frame to be received from the destination.

rogue devices Wireless devices installed on a network without authorization or verified configurations.

shared key authentication An 802.11 authentication method that uses a standard challenge-response process with shared key encryption. It relies on a wireless station having the correct key to encrypt the AP's challenge text as the criterion for acceptance. The method's main security drawback is that it encrypts the shared key using WEP, which is easily cracked.

site survey An in-depth examination of a proposed wireless network site that is designed to determine AP placement and identify stray signals, noise, and obstacles.

sniffing Capturing network traffic during transmission.

Temporal Key Integrity Protocol (TKIP) An encryption method devised as a replacement for WEP in WPA. TKIP is based on the same RC4 mechanisms used in WEP, but it includes a method for generating new keys for each packet. It also incorporates automatic AP-client synchronization of unicast encryption keys, so keys do not have to be changed manually, as they do in WEP.

wardriving Driving around with a laptop and WNIC in RF monitor mode to detect unsecured wireless signals.

Wi-Fi Protected Access (WPA) A subset of the 802.11i standard that addresses encryption and authentication; it uses IEEE 802.1x or preshared keys for authentication and uses TKIP for encryption. See also *IEEE 802.1x* and *Temporal Key Integrity Protocol (TKIP)*.

Wi-Fi Protected Access version 2 (WPA2) The encryption and authentication architecture based on the final ratified IEEE 802.11i standard. WPA2 uses preshared keys or IEEE 802.1x for authentication and uses AES for encryption. See also *Advanced Encryption Standard (AES)* and *IEEE 802.1x*.

Wi-Fi Protected Setup (WPS) A protocol from the Wi-Fi Alliance designed to automate key distribution in small office and home networks.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Review Questions

1. What function does an SSID serve on a wireless network?
   a. identifies an access point's location
   b. monitors a station's connection properties
   c. identifies the network name
   d. encrypts network traffic
2. Which of the following OSI layers is the most important in a wireless network? (Choose all that apply.)
   a. Physical
   b. Network
   c. Session
   d. Data Link
3. Which of the following frames carries TCP/IP packets in a wireless network?
   a. management
   b. data
   c. control
   d. MAC
4. What is the purpose of a beacon frame?
   a. An AP sends it to determine whether to allow a device to enter the network.
   b. It advertises services or information about the wireless network.
   c. It aids in establishing and maintaining communication.
   d. It assists in delivering frames that contain data.
5. Which of the following is the process of listening on each available channel for an AP's beacon?
   a. surfing
   b. monitoring
   c. scanning
   d. probing
6. The IEEE 802.11 standard provides which of the following authentication methods? (Choose all that apply.)
   a. asymmetric
   b. shared key
   c. open system
   d. closed system

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

7. A station can be authenticated without being associated. True or False?
8. Which is the most secure wireless implementation?
   a. WPA2-TKIP Personal
   b. WPA-TKIP Enterprise
   c. WPA2-AES Personal
   d. WPA-EAP Enterprise
9. Which authentication method does WPA2 Enterprise Security use?
   a. preshared key
   b. 802.1x
   c. EAP
   d. AES
10. Even though 802.11 wireless devices can hold up to four keys simultaneously, they have to use the same one to communicate with each other. True or False?
11. A site survey helps accomplish which of the following? (Choose all that apply.)
    a. determining placement of stations
    b. evaluating what type of antennas to use
    c. identifying where APs should be placed
    d. determining where data should be stored
12. Having an abundance of APs in a wireless network improves security. True or False?
13. Which of the following default settings should you change before connecting a device? (Choose all that apply.)
    a. SSID
    b. channel
    c. port address
    d. MAC address
14. WPA2 uses which of the following for encryption?
    a. LEAP
    b. EAP
    c. AES
    d. EAS
15. Which control frame gives a station permission to transmit?
    a. RTS
    b. CTS
    c. AES
    d. ACK

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Projects

Hands-On Project 7-1: Examining Physical Security for Wireless Devices

Time Required: 30 minutes

Objective: Research the issue of physical security for wireless devices.

Description: In this activity, you search the Internet to find examples of devices that physically secure wireless components. Select a device and write a summary of it.

1. Start your Web browser, and perform an Internet search for wireless physical security devices.
2. Visit several sites, and examine at least three different devices.
3. Select one device and write a two- to three-paragraph summary of it. Include details such as the device's cost, its purpose, how it works, and where to get it. Be sure to cite your sources.
4. Proofread your paper, and submit it to your instructor.

Hands-On Project 7-2: Finding Default Values for Wireless Devices

Time Required: 15 minutes

Objective: Research default values for popular wireless devices.

Description: In this activity, you search the Internet to find default values for popular wireless devices.

1. Start your Web browser, and open a search engine to find Web sites for the manufacturers listed in the following steps.
2. What is the default password for a Netgear wireless 802.11n router? What is the default username?

   _____

3. What is the default SSID for a Linksys wireless 802.11n router? What is the default password?

   _____

4. What is the default password for a Cisco Aironet AP? What is the default username?

   _____

5. Be prepared to discuss your findings in class. Leave your system running for the next activity.

Hands-On Project 7-3: Researching Security Products for Handheld Devices

Time Required: 30 minutes

Objective: Research ways to implement security for handheld devices.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Description: In this project, you research products that can help you implement company security policies for wireless handheld devices.

1. Review the list of best practices from the "Best Practices for Wireless Network Security" section earlier in this chapter.
2. Select three of the best practices, and then search online for products that you can use to implement these practices. Identify three different products that can enforce one best practice each.
3. Write a report that lists the three best practices you chose and the products that enforce these practices. Describe the advantages and disadvantages of each solution.
4. Submit the report to your instructor, and be prepared to discuss your findings in class.

Hands-On Project 7-4: Creating a Presentation on Wireless Attack Methods

Time Required: 2 hours

Objective: Research current wireless attacks and develop a presentation for use in security awareness training.

Description: In this project, you research current techniques for wireless attacks and develop a presentation you can use in security awareness training.

1. Start your Web browser, and search for current wireless attack techniques.
2. Select a common attack technique, and then develop a presentation to explain how it takes place, what network vulnerability the attack exploits, and what countermeasures can be used to defend against it. You can write a speech or memo, create visual aids such as charts, or develop a PowerPoint presentation.
3. Design your presentation for nontechnical users, but be sure you can answer any technical questions that might come up.
4. Submit your work to the instructor or present it in class, depending on the instructor's directions.

Hands-On Project 7-5: Observing How People Handle Wireless Devices

Time Required: 2 hours

Objective: Observe how people handle their wireless devices and note any security lapses.

Description: In this project, you observe how people handle their wireless devices, take notes on your observations, and report your findings to your classmates or instructor. You need a notepad and pencil for this project.

1. Choose a busy time to observe wireless users.
2. Select a location where you are likely to find plenty of people with wireless devices, such as your school's common areas or a cybercafé.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

3. Watch the people around you and observe how they handle their cell phones, laptops, and other wireless devices.
4. Did you notice any mishandling or any security lapses? For example, did anyone leave a cell phone unattended so that it could be stolen or tampered with? Did anyone leave a laptop unattended or unsecured? Did anyone place a device in a potentially damaging environment, such as beside a hot or cold drink?
5. Write down your observations. Turn them in to your instructor or present them in class.

Case Projects

Case Project 7-1: Planning a Wireless Network Site Survey

You are a network engineer for a small motel chain that has just been acquired by a large hotel company. The plan is to keep your company's motels as a separate brand from the new parent company. The parent company wants to upgrade the facilities in your motel chain, including the wired LANs. The company plans to implement WLANs both to assist in motel administration and to provide guests with a wireless Internet connection. You have been asked to prepare a 15-minute oral presentation on the process of developing a wireless network site survey as part of the feasibility study. The presentation should include PowerPoint slides and should explain the purpose of a site survey, the specific procedures involved, the reasons for the procedures, security considerations, and the tools typically needed to complete the survey. You have heard that the parent company intends to consolidate many departments, including information technology, so you are determined to do an excellent job on the presentation to impress your new supervisors. Prepare the presentation, and then submit your work to the instructor or present your work in class.

Case Project 7-2: Summarizing Research

When preparing to develop a new IT project, it is often necessary to perform detailed research to determine what experiences others have had with similar initiatives, what problems were encountered, and what successes can be leveraged. You learn more about security policies in Chapter 13, but in this project you summarize research that would help you prepare to develop a security policy for a mobile device.

For this assignment, read four credible research articles and then write a one- to two-paragraph summary of each article. The first two articles are provided for you. Find two additional credible articles that relate to mobile device security and summarize them as well. During the search, keep in mind that the Internet is full of information, but not all of it is credible. For example, white papers published by commercial vendors might be biased toward their own products or solutions. You might find it helpful to consult your school's reference librarian for assistance in locating appropriate sources.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Here are the first two articles for you to summarize:

- Phishing on Mobile Devices (*www.cs.berkeley.edu/~daw/papers/mobphish-w2sp11.pdf*)
- A Survey of Mobile Malware in the Wild (*www.cs.berkeley.edu/~daw/papers/mobilemal-spsm11.pdf*)

Submit your summaries of the four articles to your instructor.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 8

Intrusion Detection and Prevention Systems

After reading this chapter and completing the exercises, you will be able to:

- Identify the components of an intrusion detection and prevention system
- Describe options for implementing intrusion detection and prevention systems
- Explain the steps of intrusion detection

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

A security policy is the basis for securing an organization's network. An intrusion detection and prevention system (IDPS) adds another level of defense to firewalls, virtual private networks (VPNs), and other security devices. Like a burglar alarm, an IDPS has sensors to detect access attempts by unauthorized users and notifies you so that you can take countermeasures. Unlike a burglar alarm, however, some IDPS devices can be configured to respond in a way that actually stops attacks.

Intrusion detection involves monitoring network traffic, detecting attempts to gain unauthorized access to a system or resource, and notifying the appropriate professionals to take countermeasures. This chapter discusses IDPS components and then explains options for designing an IDPS by distributing its components at critical locations and hosts. Finally, you examine how IDPS components operate together to protect a network.

Goals of an IDPS

A network **intrusion** is an attempt to gain unauthorized access to network resources. The term *intrusion* is a polite way of referring to an attack, which is often launched with the intention of compromising the integrity and confidentiality of network data or users' privacy. An **intrusion detection and prevention system (IDPS)** consists of more than one application or hardware device and incorporates more than just detection. Intrusion detection and prevention involve three network defense functions—prevention, detection, and response. As shown in Figure 8-1, a firewall's prevention function is complemented by the IDPS, which provides detection as well as prevention; automatic reactions by the IDPS and network administrators carry out the response.

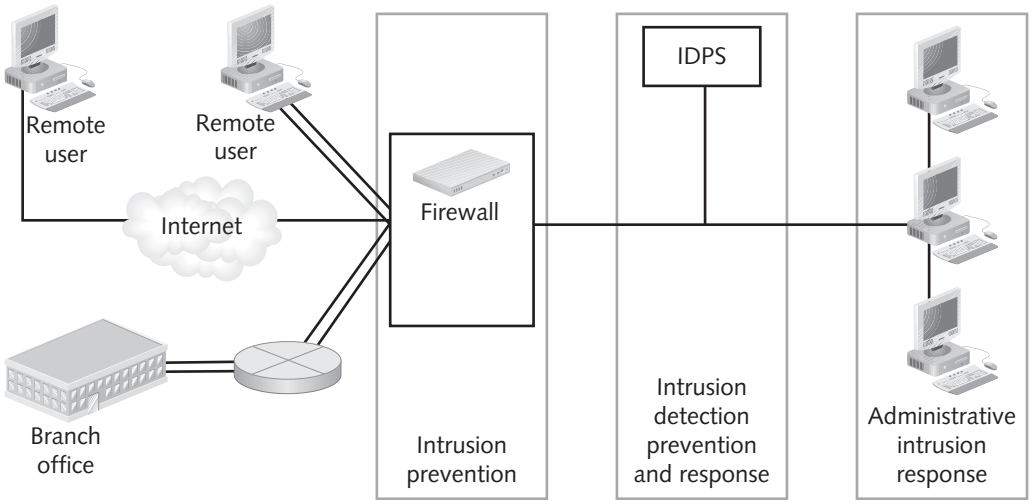

© Cengage Learning 2014

Figure 8-1 The role of intrusion detection and prevention in network defense

It may seem that the goals of an IDPS are obvious: to alert administrators to breaches of a network or individual systems, and to prevent or minimize damage by reacting to the attack

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

based on information collected by the IDPS. However, beyond the obvious goals are other considerations. It should be clear that IDPSs are only a small part of preventing intrusions. The use of firewalls, patching operating systems, and training users is also imperative to preventing network intrusions. Further, while it may not be possible to prevent an intrusion, the information gained by the IDPS might help prevent the same attack in the future.

An IDPS should be able to assess large volumes of network traffic or system activity to find signs of unauthorized access. An IDPS should also be able to record its findings in a log that allows administrators to examine past activity, and the system should be able to detect and record unauthorized access without compromise to produce evidence that will be admissible in court. Finally, an IDPS needs to be able to respond almost immediately to be effective.

An important goal of an IDPS is to make itself and the systems it protects as inaccessible as possible to attackers. With reasonable precautions, an IDPS can achieve relative security. Another value of deploying an IDPS is to demonstrate that an organization has made a good-faith effort to meet industry security standards and thus mitigate penalties associated with punitive damages if a harmful incident occurs.

Common Detection Methodologies

Anomaly and Signature Detection Systems

Currently, the three primary detection methodologies are signature detection, anomaly detection, and stateful protocol analysis. (Stateful protocol analysis, a relatively new approach, is discussed separately later in this section.) An **anomaly detection** system makes use of **profiles** that describe the services and resources each authorized user or group normally accesses on the network. Network baselines are also associated with profiles. Once these profiles are in place, the system can monitor users and groups for suspicious activity (anomalies) that does not fit the profiles. You might use anomaly detection if you are especially concerned about network misuse within the organization or you want to monitor all e-mail traffic, Web usage, and FTP servers.

Some IDPSs can create baselines during a "training period"—during this time, the IDPS monitors network traffic to observe what constitutes normal network behavior. Otherwise, you need to create profiles yourself. Because a large-scale network might consist of hundreds or thousands of users divided into many groups, profile configuration can require extensive work.

The accuracy of profiles has a direct impact on effective detection of anomalies. If profiles are accurate, the IDPS sends alarms only for genuine attacks. If profiles are incomplete or inaccurate, the IDPS sends alarms that turn out to be **false positives**—alarms generated by legitimate network traffic rather than actual attacks. False positives waste valuable time and resources; if they occur often enough, IT employees might not take alarms seriously. You also need to configure an IDPS accurately enough to avoid **false negatives**, genuine attacks that an IDPS does not detect. False negatives have the most serious security implications. **True negatives** are legitimate communications that do not set off an alarm. The term **true positive** is sometimes used to describe a genuine attack that an IDPS detects successfully.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

An anomaly-based system can also generate false positives caused by changes in user habits; after all, people do not use computer systems in the same way all the time. When users vary a pattern (by attempting to access a database they have never used before, for instance), a false positive is likely.

In contrast to anomaly-based detection, which triggers alarms based on deviations from normal network behavior, **signature detection** triggers alarms based on characteristic signatures of known external attacks. You might decide to use signature detection if you have the time and ability (and perhaps the software) to analyze the large amount of log file data this system generates.

You might also see signature detection referred to as *misuse-based detection* or as a *knowledge-based IDPS*. The names for this detection technique vary, but the method is the same: Use databases of signatures of known attacks to identify and possibly respond to intrusions.

A signature-based IDPS is good for organizations that want a basic IDPS and are mostly concerned with known attacks from intruders trying to access internal hosts from the Internet. Network engineers research well-known attacks and record rules associated with each signature; a database of these signatures is then made available to an IDPS so that it can begin protecting networks immediately after installation. Signatures should be updated regularly. An anomaly-based IDPS, on the other hand, must be trained to recognize normal network traffic before it can protect a network. Table 8-1 summarizes the advantages and disadvantages of these detection systems.

Detection method	Advantages	Disadvantages
Anomaly	Because an anomaly detection system is based on profiles an administrator creates, an attacker cannot test the IDPS beforehand and anticipate what will trigger an alarm.	Configuring the IDPS to use profiles of network users and groups requires considerable time.
	As new users and groups are created, IDPS profiles can be updated to keep up with these changes.	Updating IDPS profiles can be time consuming.
	Because an anomaly detection system does not rely on published signatures, it can detect new attacks.	The definition of what constitutes normal traffic changes constantly, and the IDPS must be reconfigured to keep up.
	The system can detect attacks from inside the network by employees or attackers who have stolen employee accounts.	After installation, the IDPS must be trained for days or weeks to recognize normal traffic.
Signature	This approach makes use of signatures of well-known attacks.	The database of signatures must be updated to maintain the IDPS's effectiveness.

© Cengage Learning 2014

Table 8-1 Advantages and disadvantages of detection systems (*continues*)

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Detection method	Advantages	Disadvantages
	This IDPS can begin working immediately after installation.	New types of attacks might not be included in the database.
	This IDPS is easy to understand and less difficult to configure than an anomaly-based system.	By making minor alterations to an attack, attackers can avoid matching a signature in the database.
	Each signature in the database is assigned a number and name so that the administrator can specify which attacks should set off an alarm.	Because a misuse-based system requires a database, extensive disk storage space might be needed.

© Cengage Learning 2014

Table 8-1 Advantages and disadvantages of detection systems (*continued*)

Monitoring an attack in progress can be helpful if you want to gather information for identifying, capturing, and prosecuting intruders or learn more about the vulnerability being exploited. The circumstances in which you allow an attack to progress should be spelled out clearly in the security policy. However, if the attack is causing severe harm, such as failure of network services or theft of proprietary data, stop the attack as quickly as possible.

Stateful Protocol Analysis

A signature-based system has another potential weakness you should keep in mind: the need to maintain **state information** (data about a connection) on a possible attack. This information gathering is called **stateful protocol analysis**. When an IDPS receives a packet, information about the connection between the host and remote computer is compared to entries in the state table. A state table maintains a record of connections between computers that includes the source IP address and port, destination IP address and port, and protocol. Furthermore, the IDPS needs to maintain state information for the entire length of the attack, which is called the **event horizon**. Maintaining this information might require an IDPS to review many packets of data; during long attacks, such as those that last from user logon to user logoff, the IDPS might not be able to maintain the state information long enough, and the attack could circumvent the system. Stateful protocol analysis can involve a variety of approaches:

- *Traffic rate monitoring*—If the IDPS detects a sudden and dramatic increase in traffic, such as that caused by a denial of service (DoS) attack, the IDPS can stop and reset all TCP traffic.
- *Protocol state tracking (stateful packet filtering)*—Some IDPSs can go a step beyond matching packet signatures by performing stateful packet filtering like that performed by firewalls. The IDPS maintains a record of the connection's state and allows packets to pass through to the internal network only if a connection has been established already.
- *Dynamic Application layer protocol analysis*—Sometimes an attacker can circumvent an IDPS by using a nonstandard port for an Application layer protocol. Dynamic Application layer protocol analysis involves detection of the protocol in use, followed by activation of analyzers that can identify applications not using standard ports.
- *IP packet reassembly*—Some IDPSs can reassemble fragmented IP packets to prevent fragments from passing through to the internal network.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Another detection method is *heuristics*—using an algorithm to detect suspicious traffic. Although this method is useful for detecting certain attacks, it is resource intensive and requires extensive tuning and maintenance to conform to the environment. Generally, heuristics are used for specific reasons and situations—for example, recognizing potential e-mail-based attacks. For most large networks, the drawbacks of heuristics often outweigh the advantages.

Examining IDPS Components

Until relatively recently, IDSs (intrusion detection systems) did not contain prevention functions. The newer all-in-one devices, IDPSs, have become more effective and less costly and are now used widely. An IDPS has the advantage of increasing interoperability, addressing security concerns more thoroughly, and centralizing management. Some IDPS products can be quite complex, however, and they require extensive knowledge to configure and maintain.

The following sections describe typical components of an IDPS:

- Network sensors or host-based agents that analyze and report activity; they are used with management servers that receive and manage information from sensors, analyze data, and identify some events
- Detection and prevention capabilities
- A command console for interfacing with the IDPS
- A database server that stores attack signatures or behaviors an IDPS uses to identify potentially suspicious traffic

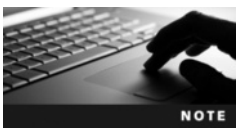

The National Institute of Standards and Technology (NIST) publishes a variety of resources, including NIST Special Publications, which are available at *http://csrc.nist.gov/publications/nistpubs.* This chapter is based on standards in NIST Special Publication 800-94, Guide to Intrusion Detection and Prevention Systems (IDPS).

Sensors and Agents

A sensor or agent functions as the electronic "eyes" of an IDPS. In host-based configurations, an IDPS installed on a single host computer has its **agent** built in to the IDPS software. In a network-based IDPS, a **sensor** is hardware or software that monitors network traffic in real time. (Host-based and network-based IDPS configurations are discussed later in the "Options for IDPSs" section.)

When a sensor detects an event it considers suspicious, an alarm is triggered. If appropriate, an automatic response may be initiated, such as blocking traffic from the source IP address of the attack. Attacks detected by an IDPS sensor can take one of two forms:

- Single-session attacks, in which an intruder makes an isolated attempt to locate a computer on the internal network or gain access by other means
- Multiple-session attacks, such as port scans or network scans, that take place over a period of time and are made up of several events

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

An IDPS that checks for network intrusions might have several sensors placed at strategic locations. Sensors should be placed at common entry points to the network (see Figure 8-2), such as the following:

- Internet gateways
- Connections between one network and another or between subnets separated by switches
- A remote access server that receives dial-up connections from remote users
- VPN devices that connect a LAN to a business partner's LAN

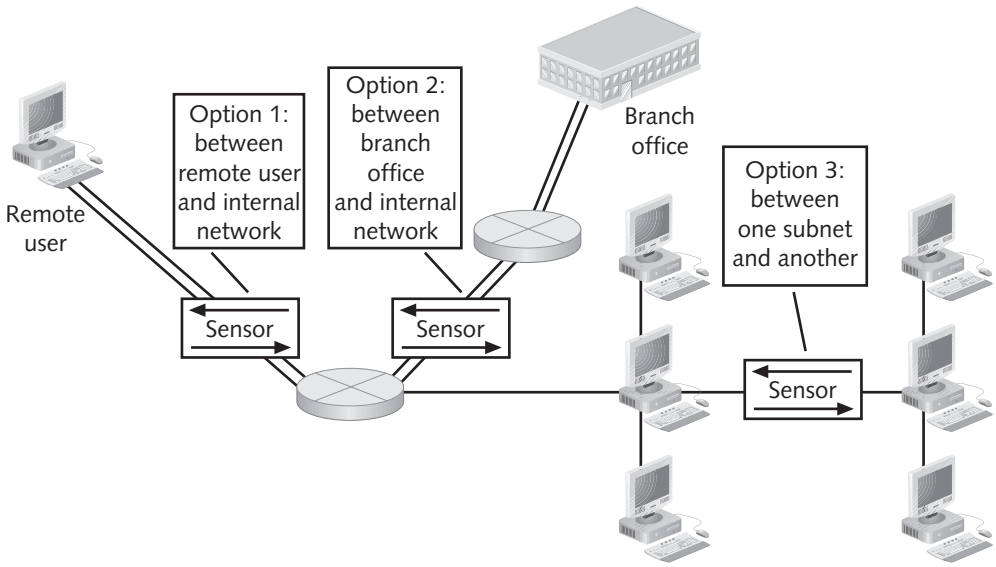

© Cengage Learning 2014

Figure 8-2 Positioning sensors at entry points to the network

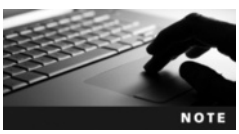

In some IDPS configurations, sensor software collects data from a hardware device called a network tap, which gathers data from network traffic traveling over the physical media.

If a firewall is used to protect the network, sensors could be positioned on either side of the firewall. However, if the sensor is placed outside the firewall at a point exposed to the Internet, it could become the subject of an attack. A more secure location is behind the firewall in the demilitarized zone (DMZ), as shown in Figure 8-3.

Management Servers An **IDPS management server** is the central repository for sensor and agent data. Sensors report log data to the management server, which analyzes and correlates the events received from several sensors. Some large networks have several management servers or even two tiers of servers, but small networks might not use any management servers.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

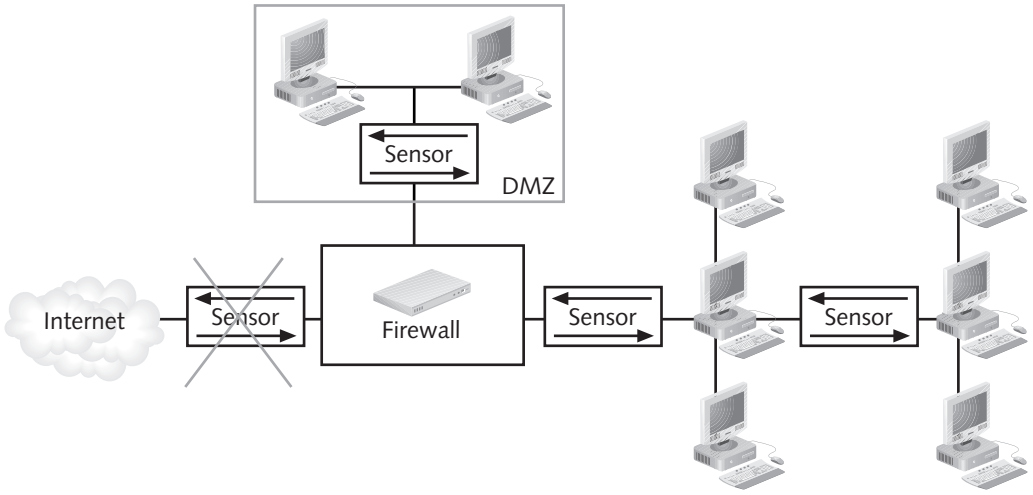

© Cengage Learning 2014

Figure 8-3 Positioning sensors behind the firewall in the DMZ

Detection and Prevention Capabilities

Most IDPSs support multiple detection capabilities for more flexibility in configuration and improved accuracy. When you are selecting an IDPS, consider the following customization options:

- *Thresholds*—Values that set the limit between normal and abnormal behavior
- *Blacklists*—Lists of entities, such as port numbers, URLs, or applications, that have been associated with malicious activity
- *Whitelists*—Lists of entities known to be harmless
- *Alert settings*—Specifying default priorities or severity levels, determining which prevention capabilities should be used for certain events, and specifying what information should be logged and how alert messages are sent, for example

Administrators can also view detection-related code, such as traffic signatures. With some IDPSs, administrators can review the application code of protocol-analysis programs. Some alerts might be generated from a complex set of signatures, so being able to modify IDPS signature code helps improve accuracy and is the only way to make the IDPS recognize characteristics that are specific to an organization. Writing custom IDPS signature code is complex and requires programming expertise because software bugs can cause an IDPS to malfunction or fail.

Any customization should be reviewed periodically to make sure it is still accurate and to account for changes in the environment; this advice also applies to baseline measurements used for anomaly detection. Keep in mind that IDPS updates and patches could affect custom settings.

Most modern IDPSs have multiple detection methods. Intrusion detection is still the basis of their operation, but these integrated systems often use compliance monitoring, too. An enterprise-class IDPS, such as McAfee Network User Behavior Analysis (*www.mcafee.com/us/products/network-uba.aspx*), monitors for vulnerabilities, access controls, logons,

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

and host-level service behaviors, and it offers custom reporting and analysis tools. Most vendors also offer integrated management consoles for organizing data collected from several systems.

Prevention Capabilities An IDPS can be configured to take preventive countermeasures, such as resetting all network connections when an intrusion is detected. Prevention capabilities vary by product. With some, administrators can specify what preventive measure should be taken for each alert type and decide whether to enable or disable prevention. Some IDPSs have a simulation mode in which all prevention capabilities are disabled, but the system generates reports that explain when the capabilities would be applied in response to different events. These reports are used to fine-tune prevention capabilities, which reduces the risk of blocking legitimate traffic.

IDPS prevention capabilities should not be considered a substitute for countermeasures taken by administrators, however. Administrators can use their judgment to determine whether an alarm is being triggered by a false positive or a genuine attack. If the attack is genuine, administrators can gauge its severity and determine whether the response should be **escalated**—increased to a higher level.

Command Console

A **command console** is software that provides an interface to an IDPS. It enables administrators to receive and analyze alert messages and manage log files. In large-scale networks with more than one IDPS, a single console enables administrators to keep up with a large volume of events so that they can respond quickly. An enterprise-grade IDPS management product usually provides a single interface for analyzing and managing security events, and may also include multiple detection and prevention capabilities.

An IDPS can collect information from security devices throughout a management network, and they are connected to the command console where they can be evaluated. When a suspicious event is detected, the command console should not be slow to respond if its host computer is busy backing up files or performing firewall functions. Therefore, a command console is usually installed on a dedicated computer to maximize response speed.

A management network can be part of an organization's regular network, but ideally it should be separate. If this arrangement is not feasible, a good alternative is a virtual management network, which you can create by setting up a virtual LAN to segregate IDPS devices. This configuration provides some protection for the IDPS, but not as much as a full management network would provide.

Database of Attack Signatures or Behaviors

An IDPS does not have the ability to use judgment, so network administrators should exercise their own judgment when evaluating security alerts. However, an IDPS can use stored information to evaluate the traffic it monitors. Signature-detection IDPSs reference a database of known attack signatures; if a sensor detects a packet or a sequence of packets that match a signature, it sends an alert. The SecurityFocus online database of known vulnerabilities (see Figure 8-4) is updated frequently; you can search it to find information about a particular type of attack *(http://online.securityfocus.com/bid)*.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

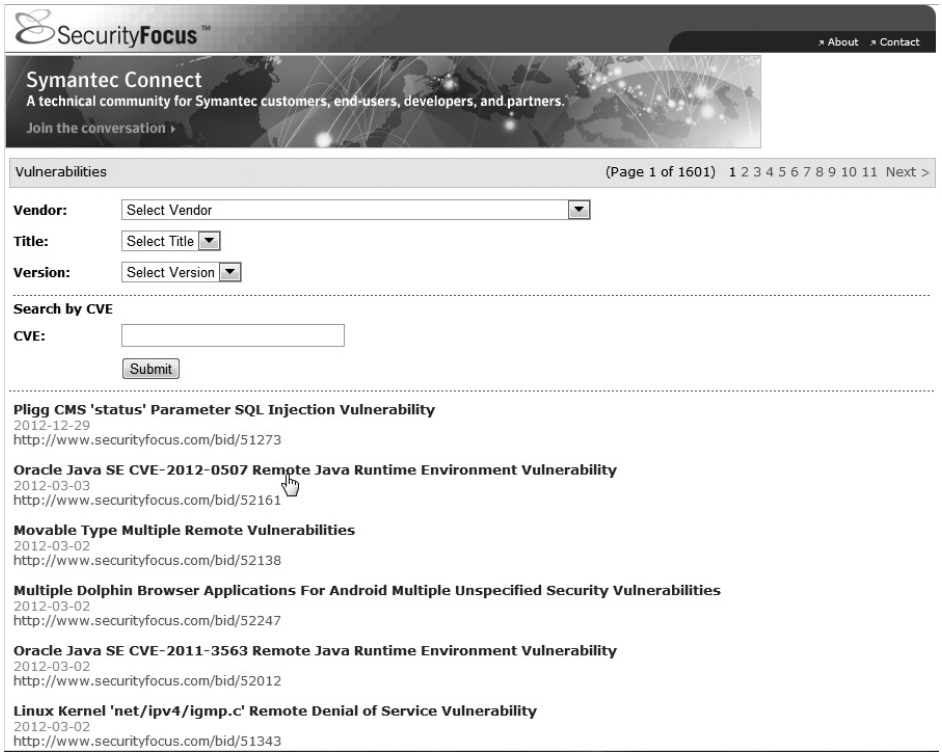

Source: Symantec

Figure 8-4 The SecurityFocus online database of known vulnerabilities

The key to using attack signature databases is to keep them up to date; a new type of attack that has not been added to the database can defeat an IDPS quickly. An IDPS vendor that uses attack signatures should include a way to download new entries and add them to the database. The problem with systems that depend solely on signatures is that they are passive: They monitor traffic, compare it to the database, and send alerts whenever a packet matches a signature, which can result in many false positives. With most IDPSs, however, administrators can address this problem by adding custom rules to the signature database. An anomaly-based IDPS also uses a database of stored information for evaluating network traffic, so custom rules can be used with these systems to reduce false alarms.

Options for IDPSs

The preceding sections described different ways that an IDPS detects suspicious events and sends alarms. The following sections examine another way to categorize an IDPS: by its position on the network and how that position affects its activities. In the following sections, you learn about network-based IDPSs, host-based IDPSs, and hybrid IDPSs.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Network-Based IDPSs

A **network-based IDPS** (**NIDPS**) is a set of IDPS components that are specialized for network use. It examines traffic on network segments by using well-positioned sensors, management servers, a command console, and databases of signatures. It also has a mechanism for storing logs, backing up configurations, taking preventive action, and sending alert messages to designated administrators.

Sensors on an NIDPS can be hardware appliances or software. An appliance-based sensor usually has specialized network interface cards (NICs) for packet capture and processing, and for analyzing traffic signatures. NIDPS sensor appliances run an OS that is specially hardened and is not usually accessed by administrators. Software-based sensors can be installed on hardened hosts or contain a specialized OS designed to run on hosts with specific configurations.

Positioning an NIDPS on the Network NIDPS sensors are commonly installed behind the firewall and before the LAN, between the firewall and the DMZ, or on any network segment (see Figure 8-5).

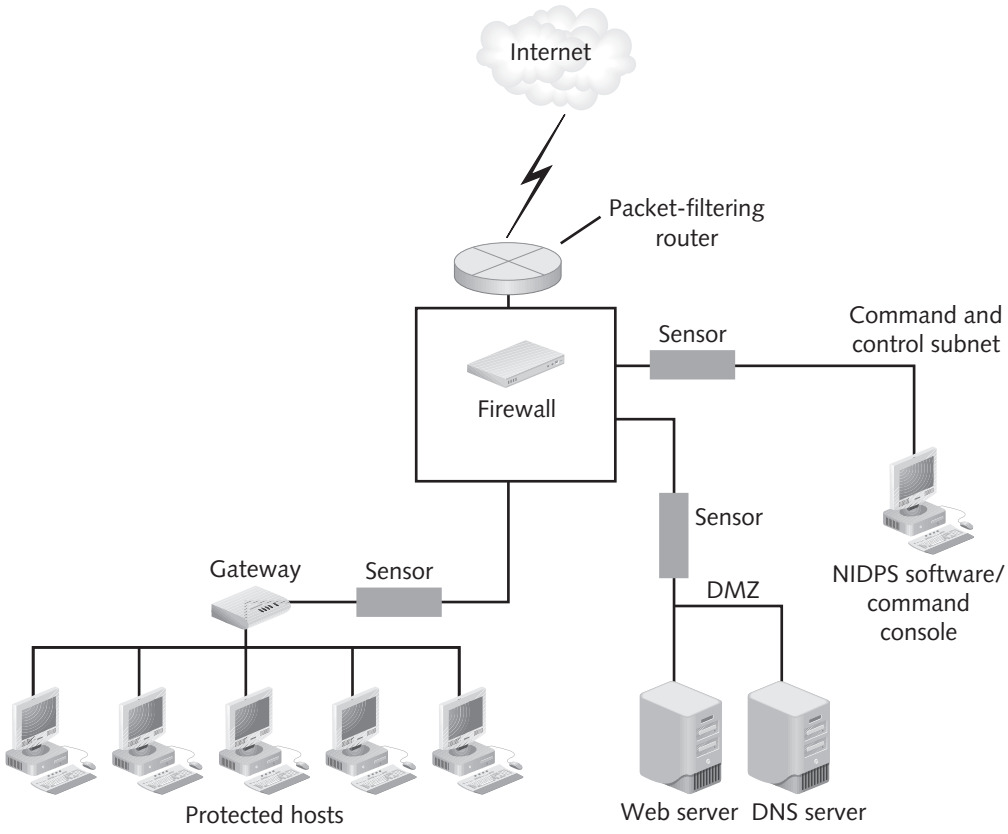

© Cengage Learning 2014

Figure 8-5 An NIDPS monitoring traffic behind the firewall or in the DMZ

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Positioning sensors at the network perimeter is ideal for enabling the IDPS to sniff packets (receive and analyze them) as they pass into the network. Each IDPS sensor is also equipped with its own NIC so that it can sniff packets in promiscuous mode, in which each packet is detected and analyzed in its entirety.

An NIDPS can use inline sensors or passive sensors. An **inline sensor** is positioned so that network traffic must pass through it. This type of sensor is used to stop attacks from blocking network traffic and is usually placed where firewalls and other security devices are positioned, such as between network segments or at connections to external networks. The drawback of inline sensors is the potential to create a traffic bottleneck if the sensor becomes overloaded, but you have two possible workarounds for this problem. You can position the inline sensor on the more secure side of the network, such as behind the firewall, so that it has less traffic to process. You can also place the sensor on the less secure side to protect and lessen the load on the device that divides the networks, such as a router. Figure 8-6 shows an inline sensor positioned inside the firewall perimeter.

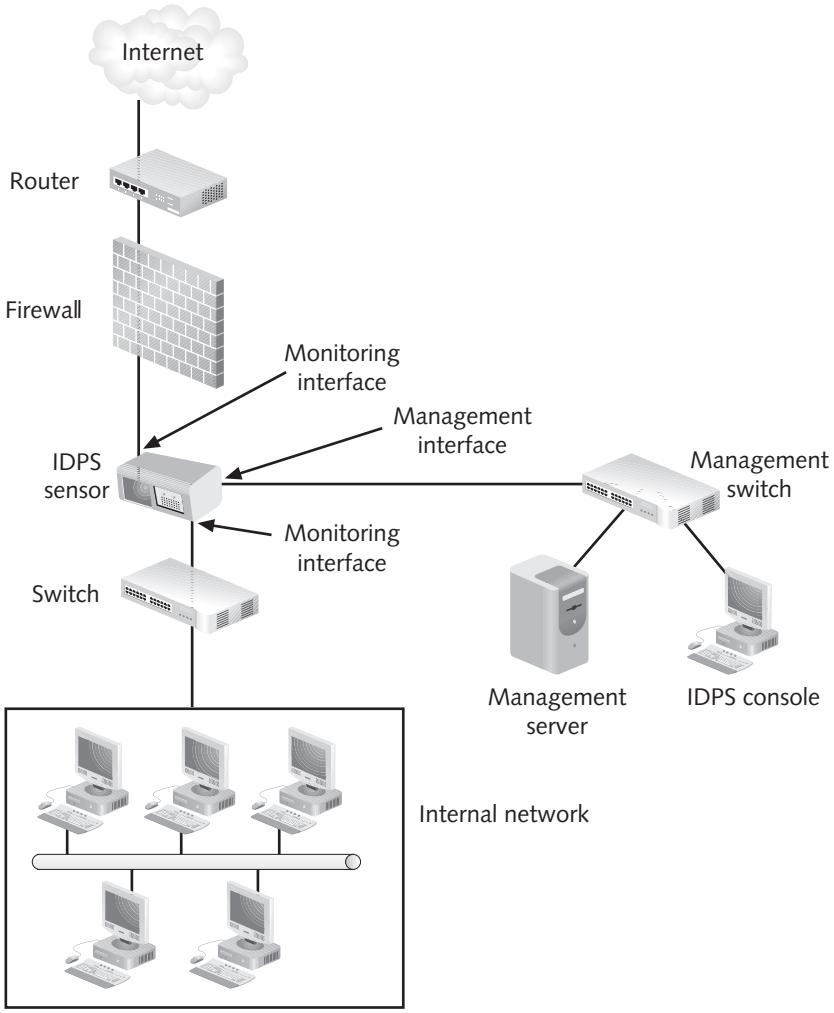

© Cengage Learning 2014

Figure 8-6 Positioning an inline sensor

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Passive sensors monitor copies of actual traffic; no actual traffic passes through them. They are typically placed at key network locations, such as divisions between networks, or on key network segments, such as the DMZ (see Figure 8-7). Passive sensors monitor traffic by using the following methods:

- *Spanning port*—A port on many switches that can see all network traffic passing through. A passive IDPS sensor attached to this port can monitor traffic flow without the traffic actually passing through it.
- *Network tap*—A direct connection between a sensor and the physical network medium, such as a fiber-optic cable.
- *IDPS load balancer*—A device that collects and directs traffic to monitoring systems. Administrators configure rules that tell the load balancer where to direct different types of traffic.

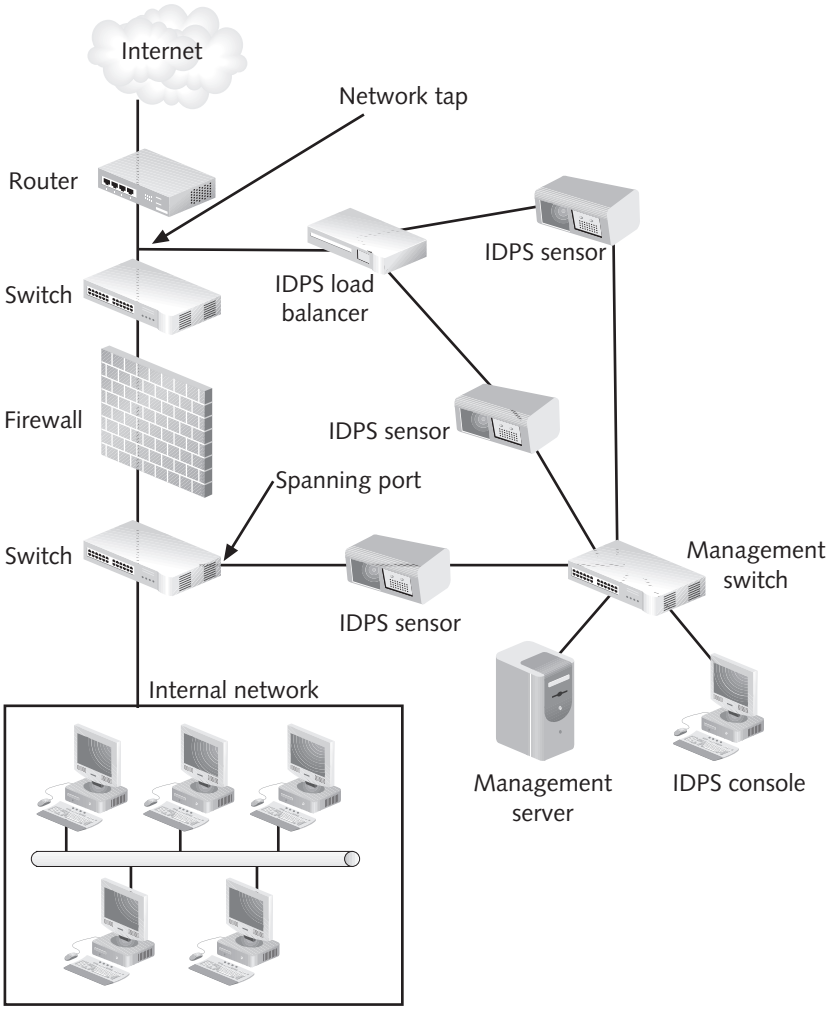

© Cengage Learning 2014

Figure 8-7 Positioning a passive sensor

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

NIDPS Capabilities NIDPS capabilities vary depending on the product. Some can collect information about hosts, OSs, applications, and network activities and characteristics. By identifying hosts, you could develop a list of all hosts on the network, organized by IP or MAC addresses. Tracking which ports are used on each host can indicate which OS is running, and knowing which OS versions are used can also help identify vulnerable hosts. An NIDPS can also analyze packet headers to identify characteristics or unusual behavior of OSs. By identifying an application, you can monitor which ports are being used and help identify misuse of the application. Collecting general information about network characteristics, such as hop count, can help identify changes to the network configuration, such as when a new device is added or removed.

A primary function of an NIDPS is extensive logging of traffic. You can use traffic logs to identify and analyze potential attacks, locate vulnerabilities, assess network use and performance, and correlate with other device logs. NIDPSs usually log the following types of information:

- Timestamps
- Event or alert types
- Protocols
- Connection or session IDs
- Source and destination IP addresses and ports
- Size of transmissions, usually in bytes
- State-related information
- Application requests and responses
- Network, Transport, and Application layer protocols
- Preventive action taken, if any

Most NIDPSs use a variety of detection capabilities, which are usually combinations of anomaly detection, signature detection, and stateful protocol analysis. These capabilities work together to improve detection efficiency and accuracy. For example, stateful inspection might parse activity into requests and responses, which are then examined for anomalies and compared to signatures. This type of multilayered inspection is similar to the defense-in-depth strategy, and is better for accuracy than using a single detection method.

The prevention capabilities of an NIDPS vary depending on the product and the sensor type used. An NIDPS has one of the following sensor types:

- *Passive only*—Ends the current TCP session (called session sniping, not commonly used now)
- *Inline only*—Uses inline firewalling and bandwidth throttling, and alters malicious content
- *Passive and inline*—Reconfigures other network security devices, perhaps by instructing a firewall to block certain types of activity or running an administrator-defined script

With most NIDPSs, administrators can configure specific actions for each type of alert. Some also offer simulation modes for fine-tuning alerts and prevention actions.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

NIDPS Management After choosing an NIDPS, an administrator should design the architecture and then test and secure the NIDPS components. Designing the architecture includes determining where sensors are located, how many are needed, and how they should be connected. Testing NIDPS components includes accounting for network downtime while deploying sensors or network taps or activating spanning ports. Securing NIDPS components involves making sure that sensors are not assigned IP addresses on monitoring interfaces so that other hosts cannot initiate connections to them. Securing components also requires hardening management networks and configuring hosts for log files and backups.

As with any networking component, NIDPS architecture, detection, and prevention configurations should be evaluated regularly. NIDPS software and hardware systems must be updated and patched periodically, and custom configurations should be monitored and adjusted as needed to remain effective against network threats.

Host-Based IDPSs

In contrast to an NIDPS on the network perimeter, a **host-based IDPS** (**HIDPS**) is deployed on hosts inside the network perimeter. On a small network, HIDPS features might be hosted on the machine that an HIDPS is monitoring. The HIDPS host could be a printer, Web server, computer, firewall, switch, router, or combination. HIDPSs are generally deployed only on sensitive or mission-critical hosts because placing one on every host in the network would be very expensive. HIDPSs in a large network environment commonly use management servers, signature databases, and consoles for configuration and monitoring. Often, an HIDPS uses the same IDPS infrastructure as an NIDPS. Figure 8-8 shows a typical HIDPS deployment.

An HIDPS can also consist of dedicated appliances running agent software that are positioned to monitor traffic on a particular host. Technically, this arrangement could be considered an inline NIDPS, but it is categorized as an HIDPS because it is highly specialized for monitoring a single host or type of traffic. An HIDPS appliance is often used to protect a Web server or database server.

An HIDPS monitors and evaluates packets generated by the host and gathers data from OS and application logs on the host. An HIDPS also gathers system variables such as the following:

- System processes
- CPU use
- File accesses
- System logs
- System and application configuration changes

System events that match the signatures of known attacks reach the IDPS on the host, which sends an alert message to users or the administrator. An HIDPS does not sniff packets as they enter the network. Instead, it monitors log file entries and user activity, and is highly effective at tracking misuse of resources by internal users.

Configuring an HIDPS An HIDPS can have two configurations: centralized or distributed. In a centralized configuration, the HIDPS sends all data it has gathered to a central location such as the command console or management server for analysis. In a distributed

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

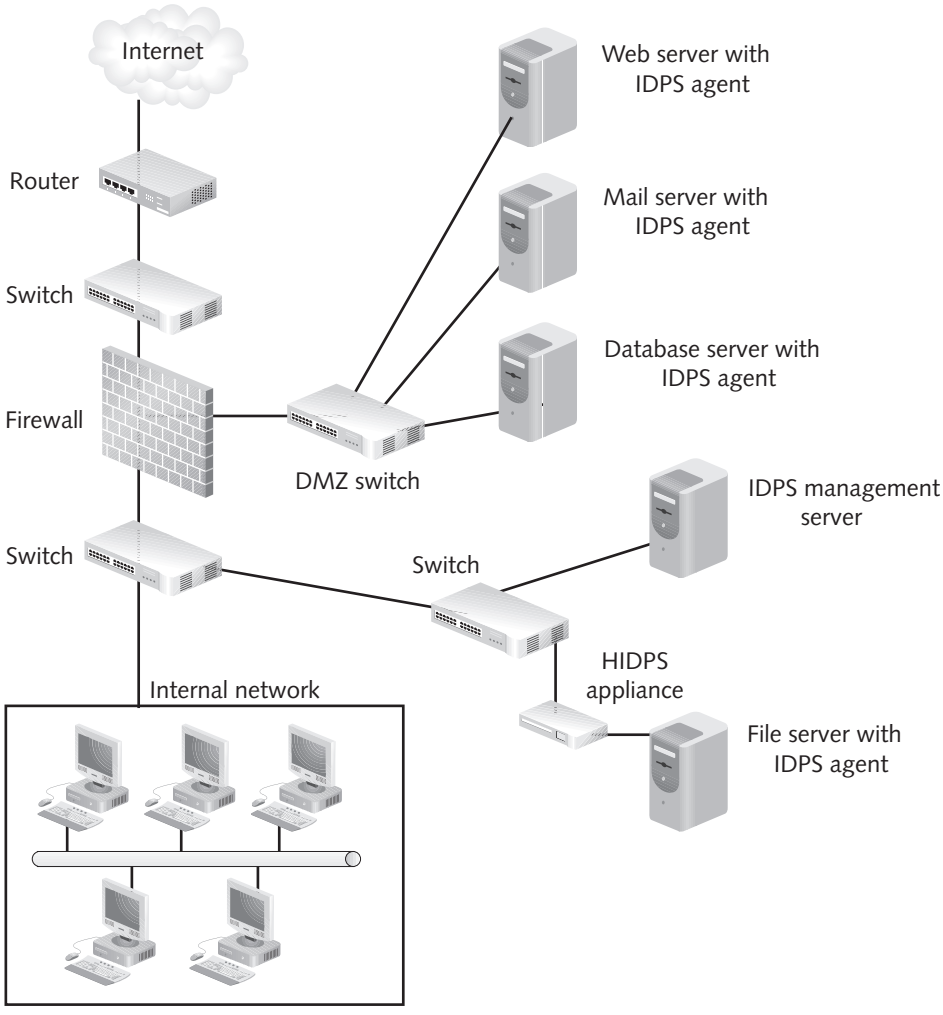

© Cengage Learning 2014

Figure 8-8 A typical HIDPS deployment

configuration, data analysis is distributed among hosts; each host analyzes data and sends only alert messages (not the data) to the command console.

In a centralized configuration (see Figure 8-9), the host's performance is unaffected by the IDPS. However, because data is sent to the command console for analysis, alert messages do not occur in real time.

The process in Figure 8-9 is as follows:

1. An event is generated on the host.
2. Data gathered by the IDPS agent (a software program running on the host) is transmitted to the command console for analysis.
3. A log file entry is created.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

4. If necessary, an alert is generated.
5. The IDPS responds.
6. Finally, data is stored in a database for long-term analysis.

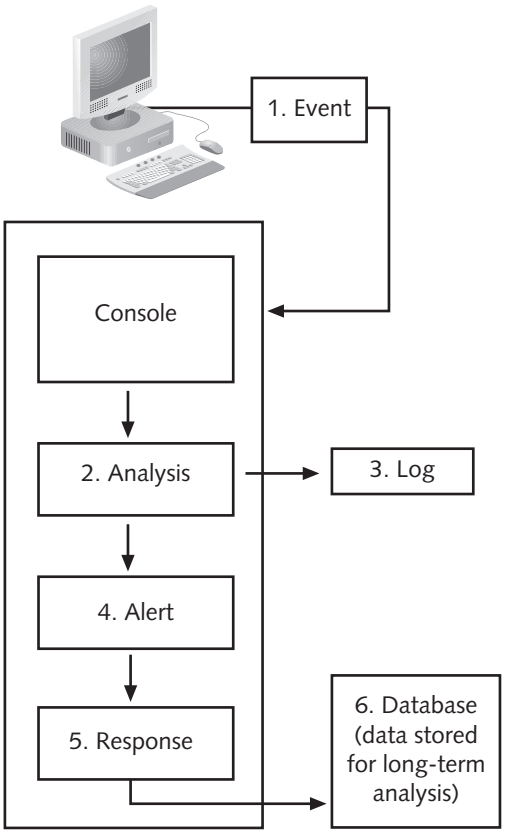

© Cengage Learning 2014

Figure 8-9 A centralized HIDPS

In a distributed configuration, event data processing is distributed between the host and command console. The host generates the data and analyzes it in real time. As a result, analysis can be performed without a delay, but the trade-off is reduced performance on the host. The host processes all data whether alerts are required or not. Data is then transmitted to the command console in the form of alert messages, as shown in Figure 8-10.

Choosing the Host The RAM, hard disk space, and processor speed required on the host depend on the type of HIDPS you use. In a centralized configuration, processing is performed on the command console, so the host's performance requirements are minimal. However, in a distributed configuration, the host gathers intrusion data and analyzes it in real time, so it needs the maximum memory and processor speed. Check the vendor's IDPS requirements for specific details.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

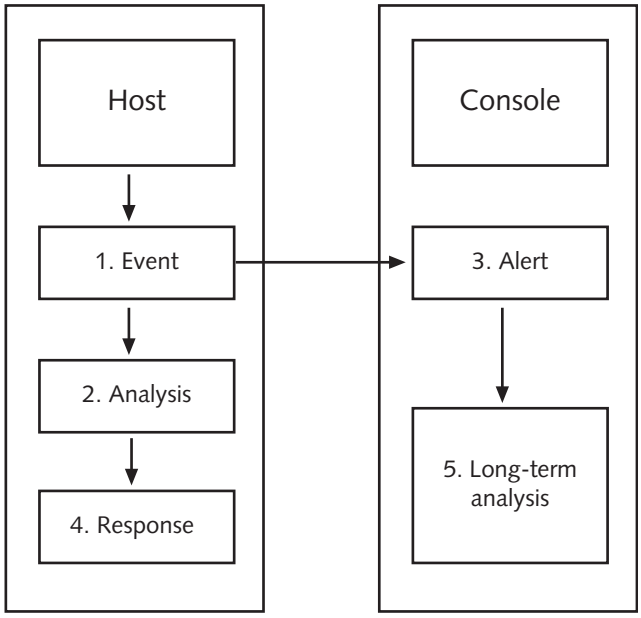

© Cengage Learning 2014

Figure 8-10 Processing event data from an HIDPS

Comparing an NIDPS and HIDPS

An HIDPS can tell you whether an attack attempt on the host was successful. An NIDPS, in contrast, provides alerts on suspicious network activity but does not tell you whether an attack attempt reached the targeted host and whether an intrusion actually occurred. An HIDPS can detect attacks that would get past an NIDPS. For example, fragmentation, out-of-sequence techniques, or other masking techniques might bypass firewalls and NIDPSs as legitimate traffic.

On the other hand, an HIDPS provides only data pertaining to the host on which it is installed, not the network as a whole. An HIDPS cannot detect an intrusion attempt that targets the entire network, such as a port scan on a range of computers. If you use an HIDPS, you need to install it on several hosts on the network, which takes time and can be more expensive than an NIDPS.

In addition, an HIDPS can compare records stored in audit logs to detect inconsistencies in how applications and systems programs are used. However, they are susceptible to some DoS attacks and could create increased performance overhead on host systems.

Hybrid IDPSs

A **hybrid IDPS** combines the capabilities of an HIDPS and NIDPS for more flexibility and security. The challenge in using a hybrid IDPS is getting the components to work together, although many IDPS products have built-in hybrid capabilities. Variations of a hybrid IDPS—combined IDPS sensor locations and combined IDPS detection methods—are described in the following sections.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Combining IDPS Sensor Locations One type of IDPS hybrid combines host-based and network-based systems. The combination enables sensors to be positioned on network segments and as agents on hosts. As a result, the network can report on attacks aimed at network segments or the network as a whole. In addition, computers that store confidential information, such as databases of job records, can be protected with an HIDPS. An IDPS on a host, especially one with a distributed configuration, can analyze data in real time and send an alert immediately that notifies the administrator of a possible unauthorized access attempt.

Combining IDPS Detection Methods Another type of IDPS hybrid combines anomaly and signature detection. The combination helps overcome the limitations of each detection method in the following ways:

- Having a database of known attack signatures enables the system to start running immediately and effectively wards off most well-known external attack methods.
- Having an anomaly-based system keeps the system flexible and capable of detecting internal misuse that deviates from normal use patterns.

A hybrid IDPS that combines anomaly and signature detection can respond to the latest, previously unreported attacks and to attacks from both external and internal sources. A drawback is that administrators must do more configuration and coordination work. Data from multiple sources must be collected in a central location where it can be reviewed and analyzed.

Advantages and Disadvantages of a Hybrid IDPS A hybrid IDPS has the advantage of combining aspects of NIDPS and HIDPS configurations. You can monitor the network as a whole with network-based sensors and monitor attacks that reach computers with host-based sensors. The drawback of a hybrid IDPS is the problem of getting disparate systems to work in a coordinated fashion. In addition, data gathered from multiple systems can be difficult to analyze easily.

Securing IDPS Components

As you might expect, IDPSs are a concern for attackers; efforts to circumvent detection are a prime focus of professional hackers. While thrill-seeking or hacktivist attackers may want to have their presence known to increase their status with other hackers or to gain attention for their cause, sophisticated criminal attackers want to remain hidden. A compromised IDPS gives hackers much more flexibility in crafting their attacks.

In some cases, a stealthy attack may not be possible, and the attacker may disable an IDPS with a denial of service attack before moving on to targets that are more valuable. One approach to address this type of attack is to create a hierarchical IDPS architecture based on mobile agent technologies so that the IDPS components can relocate from compromised hosts to intact hosts.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

IDPS Security Best Practices

- An IDPS must be able to handle the volume of traffic or activity it encounters so that it does not drop packets. Providing adequate throughput helps secure the IDPS and the network assets.
- IDPSs should be tested regularly. Testing should include security capabilities, logging, performance, and management system operation.
- Sensors should not be addressable. In other words, although sensors can capture packets, they should not have an IP address on the network they are monitoring. The addressable interface should be only on the management subnet.
- Communication between IDPS components should be encrypted.
- Authentication should be required for use and administration of the IDPS, and access control and auditing should be implemented.
- IDPSs should be able to continue operating during denial of service attacks.
- Remote logging should be used in an HIDPS.
- Operating systems of HIDPSs should be patched and hardened.

Developing IDPS Filter Rules

You have learned that an IDPS can use specific signatures of malicious packets to detect attacks. As you might imagine, these rules can get complicated. While rules to detect network scans may be relatively simple, rules to detect the preparatory steps of an attack through e-mail or the Web are much more involved. In the hands-on projects at the end of this chapter, you will download a free IDPS called Snort, and then create your own rule and test it.

Before you create your own rule, however, you need to learn the basics of Snort rule syntax. Each Snort rule has two sections: the rule *header* and the rule *options*. To examine these sections, use the following example:

alert tcp any any -> 192.168.21.0/24 111 (content:“00 01 86 a5”; msg:“mountd access”;)

The header is the opening portion in the example; the options are within the parentheses. The first field in the header is the action the system takes when a packet is detected that meets the requirements of the rule. In this case, the action field is “alert,” meaning that an alert will be logged. More sophisticated alert actions are available; for example, a message could be sent to an administrator’s desktop or e-mail account. Another action might be “log,” meaning that no alert would be registered or sent, but the contents of the packet would be recorded in the log file. The next field defines the protocol; this example specifies the TCP protocol, but other protocols can be used, like UDP or ICMP.

The next two fields define the source system’s socket—the source IP address and the source port, respectively. In the example, “any” is specified for both fields, meaning that a packet sent from any IP address and any port address would fit the rule specifications. This approach makes sense if you want to identify a type of packet regardless of whether it originates inside or outside your network. The next field is an arrow that indicates the source and destination systems; the arrow always points from the source to the destination system.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Next is the destination IP address field. In the preceding example, the specification "192.168.21.0/24" is a network address. Thus, if any host on the 192.168.21.0 network received a TCP packet from any system, this rule would apply. When specifying IP addresses, CIDR (Classless Interdomain Routing) notation is required. In CIDR notation, the IP address is followed by a forward slash and then the number of network identifier bits in the subnet mask. In the preceding example, the /24 means that the first 24 bits constitute the mask. The dotted decimal notation for /24 is 255.255.255.0. You are not limited to using a single arrow between IP addresses; you can also use the "<>" symbols to indicate that the rule applies to traffic moving in either direction.

The destination port field follows; port 111 is specified in the preceding example. This specification makes the rule much more precise because it instructs Snort to ignore any packets that are not being sent to port 111, even if all the other rule parameters are matched. This field completes the rule header section.

The options portion of the rule specifies detailed characteristics of the frame and more specific Snort actions. In the example, two options are used: "content" and "msg." The content option means that any packet containing specified content—in this case, the hexadecimal string 00 01 86 a5—will trigger the rule, assuming that the other specifications match. The second option, msg, instructs Snort to add the text that follows to the captured packet in the log. Here, any packet that meets the previous specifications will have the words *mountd access* added to the top of the packet shown in the log file, making it easier for administrators or software-auditing programs to spot the message when reviewing logs. All options must be listed within parentheses. The option name must be followed by a colon and the option must end with a semicolon, even if it is the last option listed. Snort is very particular; if your rule contains a syntax error, it will fail.

You can use an exclamation point to negate a statement. In the following example, the "!" means not to log any traffic from or to any host on the internal network (192.168.21.0/24):

log !192.168.21.0/24 any <> 192.168.21.0/24 any

Now that you have sufficient background, you can write your own rules in the hands-on projects later in this chapter.

Examining Intrusion Detection Step by Step

IDPSs operate in different ways depending on whether they are configured to react to anomalies or signatures. Despite these differences in operation, the process can be divided into general steps, shown in Figure 8-11, that apply to most IDPSs. These steps are described in the following sections.

Step 1: Installing the IDPS Database

The first step of intrusion detection occurs before any packets are detected on the network. Along with the IDPS software and hardware, you need to install the database of signatures or user profiles, which gives the IDPS a set of criteria against which it can compare packets passing through the sensor.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

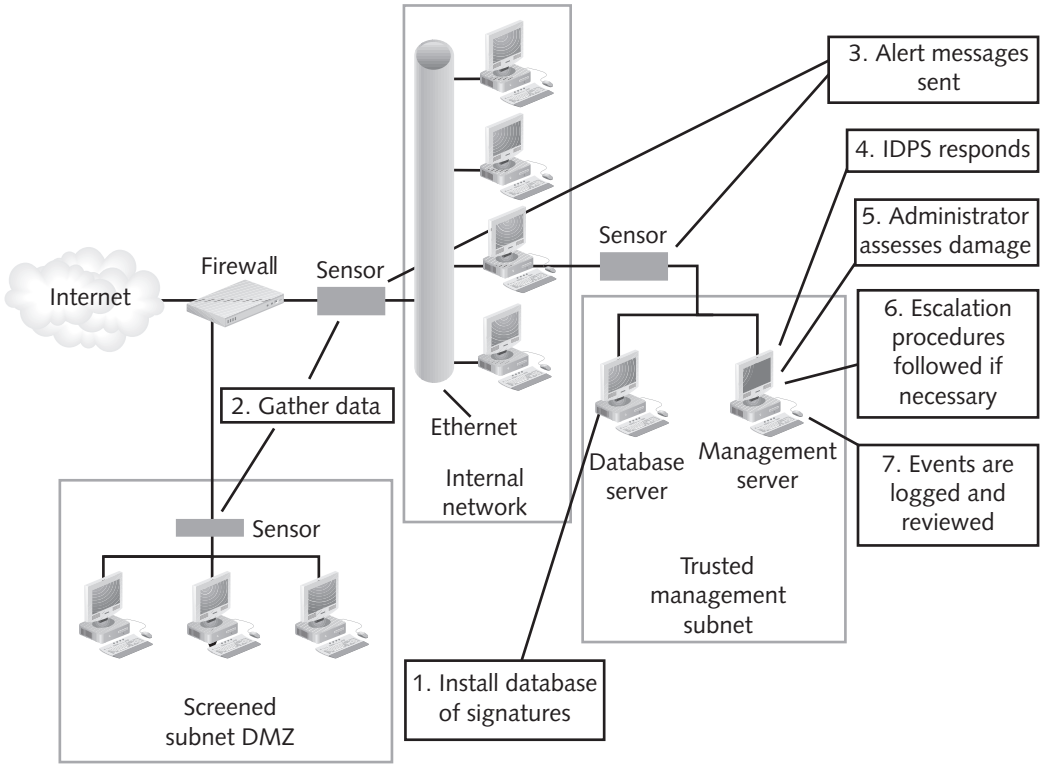

© Cengage Learning 2014

Figure 8-11 Steps in intrusion detection

In an anomaly-based system, installing the database can take up to a week longer than installing other IDPS devices so that the IDPS can observe network traffic and compile baseline data on normal network use. Some data can take a week to record because it occurs over a period of days—for example, a series of daily logons to the network. In a signature-based IDPS, you can install the database of attack signatures included with the software, or you can add your own custom rule base to account for new attacks or special situations that generate false positives.

Step 2: Gathering Data

After the IDPS and database are installed, network sensors and agents can gather data by reading packets. Agents installed on hosts observe traffic entering and leaving the hosts; sensors placed on network segments read packets that pass through those segments.

Sensors need to be positioned where they can capture all packets entering and leaving a host or network segment. Sensors placed on network segments cannot always capture every packet if the traffic level becomes too heavy, however. Repositioning agents on each network host improves accuracy, even though the expense of purchasing new agents and the effort of installing them can be considerable. The most important consideration is being able to capture all packets so that none can circumvent the IDPS.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Step 3: Sending Alert Messages

The IDPS detection software compares packets it observes with signatures stored in its database. An alert message is transmitted when a packet matches an attack signature or deviates from normal network use. The alert message goes to the IDPS command console, where the network administrator can evaluate it.

Step 4: The IDPS Responds

When the command console receives an alert message, it notifies the administrator using a method that the administrator has configured. The console might display a pop-up window or send an e-mail, for instance. An IDPS can also be configured to take action when a suspicious packet is received and an alert message is sent. The following list describes typical preventive or response actions an IDPS can take:

- *Alarm*—An alert message is sent to the command console or another designated location.
- *Drop*—The packet is dropped without an error message being sent to the originating computer.
- *Reset*—The IDPS is instructed to stop and restart network traffic, thus halting severe attacks.
- *Code analysis*—The IDPS can prevent malicious code from running or stop certain applications from opening shells used to launch some attacks.
- *File system monitoring*—The IDPS can prevent files from being modified, accessed, replaced, or deleted.
- *Network traffic filtering*—The IDPS can act as a host-based firewall to stop violations of unauthorized access or acceptable use policies, based on IP address, protocol, or port information.
- *Network traffic analysis*—The IDPS can stop incoming traffic from reaching a host or leaving it. This action is useful for stopping Network layer, Transport layer, or Application layer attacks as well as unauthorized applications and protocols.

An IDPS stops TCP traffic by sending a TCP packet with the RST (reset) flag set, which terminates the connection with the computer that is attempting to attack the system. Resetting TCP traffic does not affect UDP or other types of traffic, however. This action, called session sniping, is not used often in a modern IDPS.

Step 5: The Administrator Assesses Damage

An automated response sent by an IDPS is like a call to action. The administrator has the responsibility to monitor alerts and determine whether countermeasures need to be taken. When an IDPS is first installed, it might send many alerts that are false positives, depending on the accuracy of information in the database. Administrators usually have to fine-tune databases to account for situations that seem to be intrusions but are actually legitimate traffic.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

In an anomaly-based system, for example, an adjustment might be needed for an employee who logs on over the weekend instead of during the standard workweek. In a misuse-based system, an adjustment can be made to allow traffic that the IDPS might determine to be suspicious, such as a vulnerability scan performed by a device at a particular IP address. You could add a rule that changes the IDPS action in response to traffic from that IP address. Figure 8-12 shows the dividing line between acceptable and unacceptable network use.

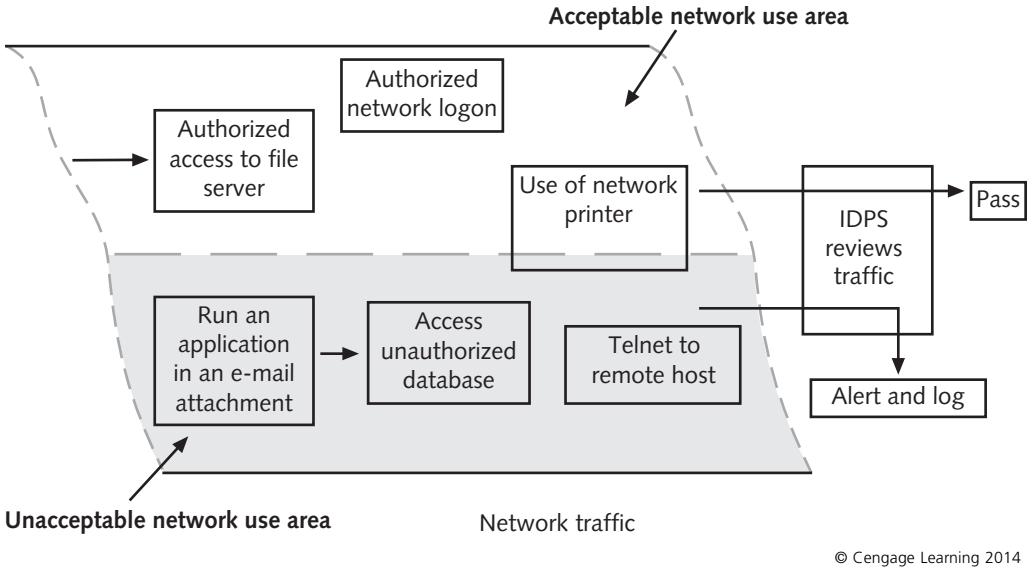

© Cengage Learning 2014

Figure 8-12 Differentiating acceptable and unacceptable network use

The line dividing acceptable network use from unacceptable use is not always clear, however. In Figure 8-12, for example, the box that indicates network printer use overlaps into the unacceptable area because acceptable use of the printer is limited to specific purposes. Printing office-related documents constitutes acceptable use, but printing personal photos probably falls on the unacceptable side of the line.

The goal of adjusting the IDPS database is not to avoid false positives because they are almost inevitable. False positives do consume an administrator's time and energy, but they do not compromise network security. The goal is avoiding false negatives—incidents that should cause an alarm but do not. False negatives occur without anyone's knowledge and are a potentially serious breach of security. Although false positives are often seen as nuisances, they are far less serious than false negatives.

Step 6: Following Escalation Procedures

Escalation procedures are a set of actions that are spelled out in the security policy and followed if the IDPS detects a true positive (a legitimate attack). These procedures vary depending on the severity of the incident. A Level One incident might be managed quickly with only a single security professional. A Level Two incident represents a more serious

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

threat and must be escalated to involve a security professional with more authority. A Level Three incident represents the highest degree of threat.

Step 7: Logging and Reviewing Events

After an IDPS has sent an alert to the command console and responded as necessary, the event that caused the alert is entered in the IDPS log. The event can also be sent to a database file, where it can be reviewed with other previous alerts. Reviewing a number of alerts sent over a period of time enables administrators to determine whether patterns of misuse have occurred. This review also gives administrators the opportunity to spot a gradual attack, such as a series of logons occurring only once every few days or a series of ping sweeps occurring once a week over a few months. As you learned in Chapter 3, a ping sweep is an attempt to solicit a response from IP addresses inside a network.

An IDPS should also provide **accountability**—the ability to track an attempted attack or intrusion back to the responsible party. Some systems have a built-in tracing feature that attempts to locate the IP address associated with an event. Identifying the person who used the source computer can be difficult, but a trace can at least provide a starting point for identifying an attacker.

Evaluating IDPS Products

Some IDPS products include antivirus or firewall capabilities, and IDPS products can be integrated and combined for comprehensive protection. The options for custom defenses that include an IDPS are numerous, so you should focus on evaluating the system that best fits your needs.

As with any network security component, the most expensive or full-featured product is not always the best option. Important considerations include the availability of staff to install and support technologies, the organization's security stance, and the operating environment. The best approach is to know what you need before you invest resources. You should consider the following basic factors when evaluating IDPS products:

- Determine whether an IDPS is necessary. Sometimes, existing network components can perform IDPS functions well enough to meet security requirements.
- Conduct a risk assessment to decide which network resources require protection and at what level.
- Define general requirements and goals that the IDPS should meet. Assess security policies and other related organizational policies that affect security and IDPS goals.
- Determine whether it is acceptable to use proprietary products or open-source products.
- Consider the frequency and accuracy of signature updates.
- Assess the availability of support.
- Evaluate the technical specifications of IT systems so that you can estimate how many products are needed and where they should be positioned.
- Determine your external security requirements, such as applicable laws, security audit requirements, results of investigated security incidents, or organizational cryptography specifications.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Evaluate your needs for security capabilities, information gathering, and logging.
- Review the detection and prevention capabilities of the IDPS products you are considering.
- Identify performance and management requirements for the IDPS.
- Define the interoperability and scalability potential for each product you are considering.
- Determine a reasonable cost estimate that includes acquisition, testing, installation, daily use, maintenance, and updates to keep the IDPS operating at peak efficiency.
- Identify resource limitations, such as budget, staff, or hardware.
- Identify any training, documentation, and support required for IDPS products.

Chapter Summary

- Intrusion detection and prevention systems (IDPSs) add another line of defense behind firewalls and antivirus software. Many IDPSs can go a step beyond transmitting alarms and actually respond to an event, which makes them more useful than older intrusion prevention systems (IPSs).
- IDPS components include sensors or agents, management servers that gather and analyze data from sensors and agents, command consoles that provide an interface for administrators, and databases of attack signatures for misuse detection or user profiles for anomaly detection.
- A network-based IDPS (NIDPS) uses sensors positioned at key points on the network. A host-based IDPS (HIDPS) deploys agents on selected hosts in the network. A hybrid IDPS combines aspects of network-based and host-based IDPS configurations. Many IDPS products can also combine detection and prevention capabilities.
- Selecting an IDPS requires evaluating the organization's needs and security goals and the product's features.
- The steps of intrusion detection typically follow this process: installing the IDPS and database of attack signatures or user profiles, gathering data, sending an alert, responding to the alert, assessing damage, following escalation procedures, and logging and reviewing events. Specified actions, such as alerts, are taken when a packet matches an attack signature or deviates from normal network use.

Key Terms

accountability The ability to track an attempted attack or intrusion back to its source.

agent An IDPS component that monitors traffic on a specific host.

anomaly detection A type of detection system that sends an alarm when it detects an event that deviates from behavior defined as normal in the IDPS profile.

command console Software that provides a graphical interface to an IDPS.

escalated Increasing an intrusion response to a higher level.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

event horizon The entire length of an attack, from the first packet the IDPS receives to the last packet needed to complete the attack.

false negatives Attacks that are not detected by an IDPS.

false positives Alarms generated by legitimate network traffic rather than actual attacks.

host-based IDPS (HIDPS) An IDPS deployed on hosts inside the network perimeter.

hybrid IDPS An IDPS that combines the capabilities of an HIDPS and NIDPS for more flexibility and security.

IDPS management server An IDPS component that serves as the central repository for sensor and agent data.

inline sensor An NIDPS sensor positioned so that all traffic on the network segment is examined as it passes through. See also *network-based IDPS (NIDPS)*.

intrusion An attempt to gain unauthorized access to network resources and to compromise the integrity and confidentiality of network data or users' privacy.

intrusion detection and prevention system (IDPS) A network security measure that can consist of applications and hardware devices deployed on the network, hosts, or both to prevent, detect, and respond to traffic interpreted as an intrusion.

network-based IDPS (NIDPS) A set of IDPS components specialized for network use. An NIDPS examines traffic on network segments by using well-positioned sensors, management servers, a command console, and databases of signatures.

passive sensor An NIDPS sensor that examines copies of traffic on the network; no actual traffic passes through the sensor. See also *network-based IDPS (NIDPS)*.

profiles Sets of characteristics that describe network services and resources a user or group normally accesses.

sensor An IDPS component that monitors traffic on a network segment.

signature detection The process of monitoring network traffic to detect unauthorized access attempts and sending notifications so that countermeasures can be taken.

state information Data about a network connection that is typically kept in a state table.

stateful protocol analysis The process of maintaining a table of current connections so that abnormal traffic can be identified based on previous packets associated with the same transmission.

true negatives Legitimate communications that do not cause an IDPS to set off an alarm.

true positive A genuine attack detected successfully by an IDPS, in contrast to a true negative or false positive.

Review Questions

1. How can data gained from intrusion detection improve network security? (Choose all that apply.)
   a. It can help prevent future attacks.
   b. It can route traffic more efficiently.
   c. It can shield IP addresses on the network.
   d. It can help determine how to respond to security incidents.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

2. Which of the following IDPS results is the cause of greatest concern?
   a. true positives
   b. true negatives
   c. false positives
   d. false negatives
3. An IDPS management server performs which of the following functions?
   a. monitoring inbound and outbound traffic
   b. modifying state information
   c. storing and analyzing sensor data
   d. providing an IDPS interface
4. Which of the following is an IDPS detection capability you can customize? (Choose all that apply.)
   a. blacklists
   b. signatures
   c. state tables
   d. thresholds
5. Misuse-based detection is based on which feature of network traffic?
   a. user profiles
   b. normal traffic
   c. signatures
   d. user accounts
6. An anomaly-based IDPS can be circumvented in which of the following ways?
   a. new attacks
   b. changes in user habits
   c. changes in published signatures
   d. minor changes in attack methods that do not match known signatures
7. A signature-detection IDPS can be circumvented in which of the following ways?
   a. changes in attack methods
   b. a stolen user account
   c. making traffic appear normal
   d. attacks made during the IDPS training period

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

8. Which intrusion detection method can begin protecting a network immediately after installation?

9. Which intrusion detection method is almost impossible for intruders to test before attempting an attack?

10. Which IDPS activity could detect a DoS attack?
    a. protocol state tracking
    b. signature detection
    c. traffic monitoring
    d. IP packet reassembly

11. What type of IDPS should you use if your main concern is preventing known attacks?
    a. signature-based IDPS
    b. network-based IDPS
    c. host-based IDPS
    d. anomaly-based IDPS

12. An HIDPS can detect an intrusion attempt that targets the entire network, such as a port scan on a range of computers. True or False?

13. What preventive responses can an IDPS make to a possible attack? (Choose all that apply.)
    a. prevent malicious code from running
    b. drop the suspicious packet
    c. allow, reset, alarm
    d. reset all network connections

14. Which of the following is almost inevitable and should be expected after installing an IDPS? (Choose all that apply.)
    a. false negatives
    b. huge log files
    c. signatures that become outdated
    d. false positives

15. How can an inline sensor be positioned to reduce the load on a perimeter security device, such as a firewall or router?
    a. Place the sensor inside the security perimeter.
    b. Replace the firewall or router with the sensor.
    c. Position the sensor outside the security perimeter.
    d. Add another firewall or router.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Projects

Hands-On Project 8-1: Installing Ubuntu Linux

Time Required: 60 minutes

Objective: Install Ubuntu Linux.

Description: In preparation for installing an IDPS, you install Ubuntu Linux, a free, open-source operating system favored by hackers and security professionals alike. First, you download the Ubuntu installation .iso file. (Alternatively, if you are using virtualization software, you can boot to the .iso file itself to install Ubuntu.)

1. Log on to a computer that has CD-burning software installed. If your instructor prefers to issue installation CDs, skip to Step 6. Open a Web browser, and go to *www.ubuntu.com/download/ubuntu/download*. You can download the latest version, but to download the version used in this book (11.10), scroll down and click **Take a look at a full list of our previous versions and alternative downloads**. Under Select the nearest mirror, click **United States**, click **11.10**, click **release**, scroll down, and click **ubuntu-11.10-dvd-i386.iso** from the list of images.
2. Click **Save** to save the Ubuntu .iso file to your computer.
3. Use CD-burning software to create the installation CD from the .iso file.
4. Place the CD in the CD-ROM drive, and boot the computer to the CD.
5. In the Welcome window, click **Install Ubuntu**.
6. In the next window, click the **Download updates while installing** check box, and click **Continue**.
7. In the Installation type window, select **Erase disk and install Ubuntu**, and click **Continue**.
8. In the Erase disk and install Ubuntu window, click **Install Now**.
9. In the Where are you? window, select the location closest to you, and click **Continue**.
10. In the Keyboard layout window, click **Continue**.
11. In the Who are you? window, type your full name in the Your name text box, and type **Team*x*Linux** in the Your computer's name text box, where *x* is the number assigned by your instructor. Your first name should appear in lowercase in the Pick a user name text box. This is your logon name. Type **Pa$$word** in the Choose a password text box and the Confirm your password text box, and click **Continue**.
12. In the Installation Complete window, click **Restart Now**.
13. In the logon screen, log on to your account with the password **Pa$$word**.
14. Click the **Update Manager** icon (see Figure 8-13). In the Update Manager window, click **Install Updates**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

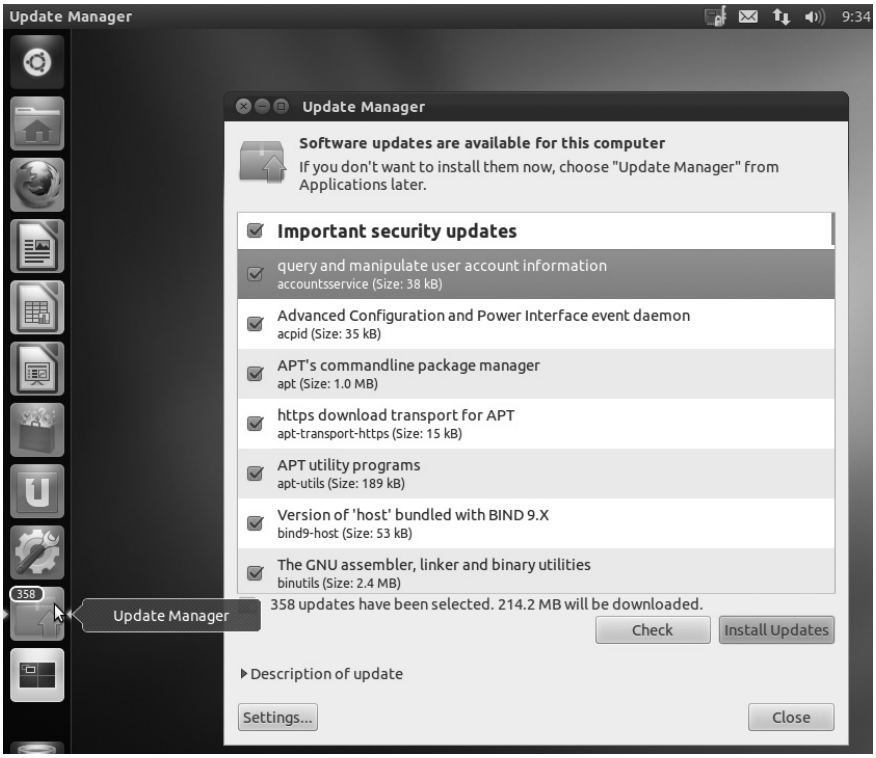

Source: Ubuntu

Figure 8-13 Ubuntu Update Manager

15. In the Authenticate window, enter your password in the Password text box, and click **Authenticate**. (Do this whenever the Authenticate window appears.)
16. Once the updates are finished installing, you can continue with the next project.

Hands-On Project 8-2: Installing an IDPS: Snort

Time Required: 10 minutes

Objective: Install Snort.

Description: In this project, you use the Ubuntu Software Center to install Snort, the most popular open-source IDPS.

1. If necessary, log on to your Linux system.
2. Roll your mouse over the icons on the left side of the desktop, find the Ubuntu Software Center icon, and click it to start the Ubuntu Software Center. In the search box in the upper-right corner, type **snort**. In the frame below, the **flexible Network Intrusion Detection System snort** appears. Select this entry and click **Install**. Enter your password in the Authenticate window, and click **Authenticate**.
3. The next window allows you to configure the address range for the local network. Enter the appropriate network address in CIDR notation—for example,

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

192.168.1.0/24. Your instructor may provide you with a specific network to use. Click **Forward**.

4. When the installations are complete, close the Ubuntu Software Center by clicking the **X** in the upper-left corner of the window.
5. In the Update Manager window, click **Restart Now**.
6. Log on to your Linux system again. Roll your mouse over the icons on the left side of the desktop, find the **Dash home** icon, and click it. In the search box at the top of the next window, type **terminal**. In the frame below, click the **Terminal** icon. Notice that the prompt in the terminal window lists the name of the user, followed by the @ sign and the name of the computer.
7. Type **snort** and press **Enter**. The next window explains the various Snort options and ends with, "Uh, you need to tell me to do something…." It is clear that you have left Windows and entered the world of open-source software.

Hands-On Project 8-3: Exploring Snort's Packet Capture Function

Time Required: 15 minutes

Objective: Capture and analyze packets with Snort.

Description: Snort has several functions. It is a packet sniffer, a packet logger, and an IDPS. In this project, you examine Snort's packet sniffing function and analyze a packet capture of traffic between Linux and Windows 7. You should work with a partner in this project.

1. If necessary, log on to your Linux system. If a terminal window is not open, follow Step 6 in Hands-On Project 8-2 to open one.
2. Explore the Snort installation files. Type **cd /** and press **Enter** to enter the root directory of the file system. Type **ls** and press **Enter** to list the contents of the / directory. Many important configuration files are stored in the /etc directory. The Snort directory is in /etc, so type **cd /etc/snort** and press **Enter**.

Be sure to observe case sensitivity in Linux. All commands are case sensitive.

3. Type **ls** and press **Enter** to list the contents of the /snort directory. Your results should be similar to those in Figure 8-14.

```
root@Team1Linux:/etc/snort# ls
classification.config  gen-msg.map       sid-msg.map  snort.debian.conf
community-sid-msg.map  reference.config  snort.conf   threshold.conf
database.conf          rules             snort.conf~  unicode.map
```

Source: Snort/Sourcefire, Inc.

Figure 8-14 The snort directory

4. The file /etc/snort/snort.conf is the main configuration file for Snort. In this file, you can specify which networks are considered internal and external, where the IDPS

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

rules are kept, how to use preprocessors, and other options. Type **cd rules** and press **Enter** to change to the rules directory. Type **ls** and press **Enter** to view the different rules that are included with Snort by default. You will create a rule in a later hands-on project in this chapter.

5. The other student partner should log on to Windows 7 with an administrative account, open a command prompt, and use the ipconfig command to verify the IP address of the Windows 7 system. From the Linux terminal window, type **ifconfig** and press **Enter** to determine the Linux IP address. It is listed on the second line as the inet addr. Record the IP addresses here:

   Linux _____

   Windows 7 _____

6. On Linux, type **sudo snort –v** and press **Enter**. You are asked for your password, but do not enter it yet.
7. At the Windows 7 command prompt, type **ping** followed by the IPv4 address of your Linux system, but do not press Enter. Ubuntu requires that you run Snort as an administrator to capture packets. The sudo command allows you to specify the administrator's account. The password is Pa$$word.
8. When both partners are ready, the Linux partner enters the correct password and presses **Enter** to start Snort. Then the Windows 7 partner presses **Enter** to start the ping.
9. When the ping is complete, the Linux partner presses **Ctrl+C** to stop Snort. A summary of the captured packets appears. Scroll up to view the captured packets; your results should be similar to those in Figure 8-15. Letters have been added to the capture to help identify eight packets. Examine the capture in Figure 8-15, and complete Table 8-2. Note that the IP address of the Linux system is 192.168.1.110 and that the Windows 7 IP address is 192.168.1.129. Figure 8-16 explains some of the fields in a Snort capture.

```
Commencing packet processing (pid=2576)
03/10-10:18:51.315092 ARP who-has 192.168.1.110 tell 192.168.1.129                A
03/10-10:18:51.315327 ARP reply 192.168.1.110 is-at 8:0:27:3E:F:1F                B
03/10-10:18:51.315610 192.168.1.129 -> 192.168.1.110                              C
ICMP TTL:128 TOS:0x0 ID:13113 IpLen:20 DgmLen:60
Type:8  Code:0  ID:1   Seq:5  ECHO
=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
03/10-10:18:51.315764 ARP who-has 192.168.1.129 tell 192.168.1.110                D
03/10-10:18:51.316829 ARP reply 192.168.1.129 is-at 8:0:27:17:25:16               E
03/10-10:18:51.316838 192.168.1.110 -> 192.168.1.129                              F
ICMP TTL:128 TOS:0x0 ID:3350 IpLen:20 DgmLen:60
Type:0  Code:0  ID:1  Seq:5  ECHO REPLY
=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
03/10-10:18:52.326261 192.168.1.129 -> 192.168.1.110                              G
ICMP TTL:128 TOS:0x0 ID:13114 IpLen:20 DgmLen:60
Type:8  Code:0  ID:1   Seq:6  ECHO
=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
03/10-10:18:52.326422 192.168.1.110 -> 192.168.1.129                              H
ICMP TTL:128 TOS:0x0 ID:3351 IpLen:20 DgmLen:60
Type:0  Code:0  ID:1  Seq:6  ECHO REPLY
=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
```

Source: Snort/Sourcefire, Inc.

Figure 8-15 Snort packet capture

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

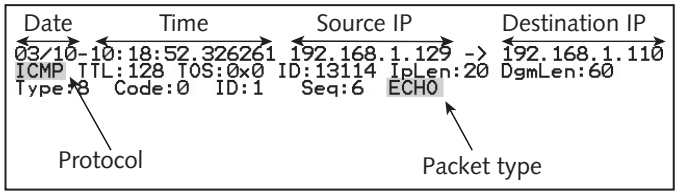

Source: Snort/Sourcefire, Inc.

Figure 8-16 Snort capture fields

Packet	Source (Windows 7 or Server)	Destination (Windows 7 or Server)	Protocol	Purpose
A				
B	Linux	Windows 7	ARP	Inform the Windows 7 system of the Linux system's MAC address
C				
D	Linux	Windows 7	ARP	Determine the MAC address of the Windows 7 system
E				
F				
G				
H	Linux	Windows 7	ICMP	The Linux system's Echo Reply in response to the second Echo Request from the Windows 7 system

© Cengage Learning 2014

Table 8-2 Packet analysis

10. Leave your systems running for the next project.

Hands-On Project 8-4: Exploring Snort's Logging Function

Time Required: 25 minutes

Objective: Capture and log packets with Snort.

Description: An important part of network forensics is maintaining logs of network activity so that you can study the traffic after an attack to help determine what systems may have been affected. Also, logs can provide valuable evidence in internal, civil, or criminal proceedings. You should work with a partner in this project.

1. On Linux, open the terminal window and type **sudo snort –vd**, but do not press Enter.
2. Open a Web browser. In the address box, type **bach-newman.com**, but do not press Enter.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

3. Open the terminal window, and press **Enter**. Immediately open the Web browser, and press **Enter**.
4. You should expect a lot of traffic. Press **Ctrl+C** quickly, and then examine the results. Figure 8-17 shows a typical packet. How is this packet different from the ones you captured in Hands-On Project 8-3? Note that data is shown in both hexadecimal and ASCII (American Standard Code for Information Interchange) formats.

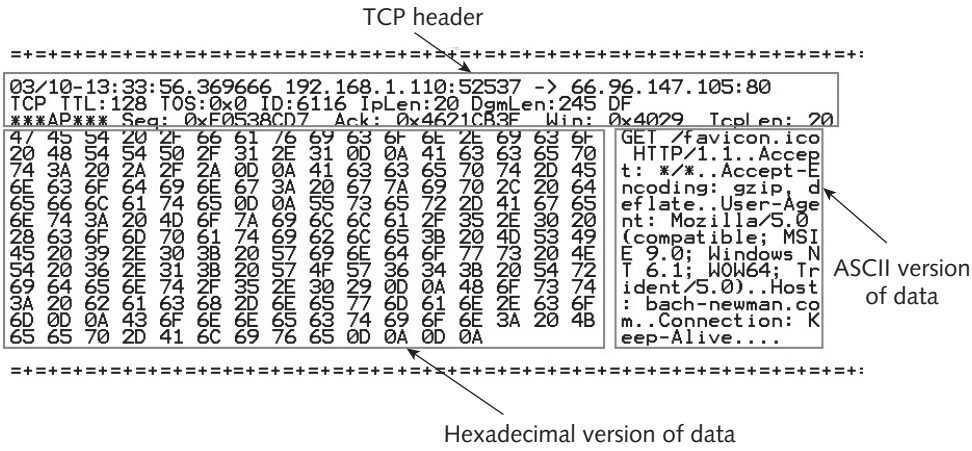

Source: Snort/Sourcefire, Inc.

Figure 8-17 Snort packet capture with the –vd option

5. To use Snort's logging function, type **cd /etc/snort** and press **Enter**.
6. Type **sudo mkdir log** and press **Enter** to create a directory called log in the /etc/snort directory. Type **ls** and press **Enter** to see the new directory listed. Note that directories are displayed in blue in Ubuntu.
7. Type **sudo snort –dv –l /etc/snort/log –K ascii**, but do not press Enter.
8. On the Windows 7 system, open a command prompt and ping the Linux system just after your partner on the server presses **Enter**.
9. On the Linux system, press **Ctrl+C** to stop the capture when the ping is complete. Type **cd /etc/snort/log** and press **Enter**. Type **ls** and press **Enter**. Note that Snort creates a directory named for the IP address of each system from which it captured packets.
10. Type **cd *IP_address_Linux***, where *IP_address_Linux* is the IP address of your Linux system. Press **Enter**. The IP address was 192.168.1.110 in Hands-On Project 8-3. Permission is denied. Your user does not have permission to enter these directories, so you must assign a password to the user named *root* who is the Linux system administrator. Then you will perform as the root user.
11. Type **sudo passwd root** and press **Enter**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

12. At the prompt, type **Pa$$word** and press **Enter**.
13. Enter the same password at the next prompt, and press **Enter**.
14. Type **su root** and press **Enter**. Type the root user's password that you entered in Step 12, and press **Enter**. Now you are acting as the root user.
15. Type **cd *IP_address_Linux*** and press **Enter**. The command should be successful this time. Type **ls** and press **Enter**. You should see a file named ICMP_ECHO_REPLY. To view the contents of the file, enter **cat ICMP_ECHO_REPLY** and press **Enter**. How are these ping packets different from the ping packets you captured in Hands-On Project 8-3?
16. Note that using the –d option in the Snort command adds the data portion of the packets. In this case, a ping packet's data includes part of the alphabet. Note also that Snort's logging feature creates directories for IP addresses and names specific log files after their function.
17. If you have time, explore the rest of the log directory. Use the **cd ..** command to go up one directory, and use the **cat** command to view files.
18. Leave your systems running for the next project.

Hands-On Project 8-5: Exploring Snort's IDPS Function

Time Required: 60 minutes

Objective: Create and test an IDPS filtering rule.

Description: In this activity, you install Samba so that the Windows 7 and Linux systems can share files. Samba is a utility that allows the Linux file and print sharing protocol, Network File System (NFS), to interoperate with the Windows file and print sharing protocol, Server Message Blocks (SMB).

Although the term *SMB* is still in use, the current Windows file and print protocol is Common Internet File System (CIFS).

Next, you create a Snort rule to detect packets that may contain a password. Then, you analyze the log files to determine if your rule was successful. You should work with a partner in this project.

1. If necessary, log on to Linux.
2. If necessary, open the terminal window. Type **su root** and press **Enter**. Type **Pa$$word** and press **Enter**. Type **apt-get install samba smbfs** and press **Enter**.
3. When asked if you want to continue, type **y** and press **Enter**.
4. Make a backup copy of the Samba configuration file by typing **cp /etc/samba/smb.conf /etc/samba/smb.conf.bak** and press **Enter**.
5. Type **vi /etc/samba/smb.conf** and press **Enter**. You are now in the Vi editor. Use the down arrow to move the cursor to the line just below the following line: # security = user. Press **i** to enter Insert mode, where you can make changes in the file. Press **Enter** once to place an additional blank line under # security = user.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

6. Type **[homes]** and press **Enter**.
7. Type **comment = Home Directories** and press **Enter**.
8. Type **browseable = no** and press **Enter**.
9. Type **writeable = yes** and press **Enter**.
10. Press **Enter** to create a blank line.
11. Type **[share]** and press **Enter**.
12. Type **comment = Team*x*Linux File Server**, where *x* is the number assigned by your instructor. Press **Enter**.
13. Type **force user = samba** and press **Enter**.
14. Type **force group = samba** and press **Enter**.
15. Type **read only = no** and press **Enter**.
16. Type **hosts allow =** and press **Enter**.
17. Press **Enter** to create a blank line.
18. Press **Esc**, type **:wq** and press **Enter** to exit the Vi editor. Your result should look similar to Figure 8-18.


```
# /usr/share/doc/samba-doc/htmldocs/Samba3-HOWTO/ServerType.html
# in the samba-doc package for details.
#   security = user

[homes]
comment = Home Directories
browseable = no
writeable = yes

[share]
comment = Team1Linux File Server
force user = samba
force group = samba
read only = no
hosts allow =

writab
# You may wish to use password encryption.  See the section on
```

Source: Samba

Figure 8-18 Samba.conf

19. Now you will add a user to the Linux system. Type **adduser linux*x*user**, where the second *x* is a number assigned by your instructor. Press **Enter**.
20. At the password prompt, type **Pa$$word** and press **Enter**. At the next prompt, type the password again and press **Enter**.
21. At the Full Name prompt, type **Linux*x***, where the second *x* is a number assigned by your instructor. Press **Enter**.
22. Do not enter anything at the next prompts, but press **Enter** until you reach the "Is this information correct?" prompt. Type **y** and press **Enter**.
23. To create a Samba password for the user, type **smbpasswd –a linux*x*user** and press **Enter**.
24. Enter the password **Pa$$word** and press **Enter**. Retype the password and press **Enter**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

25. On the desktop, click the **Settings** icon in the upper-right corner. Click **Shut Down**. In the Shut Down window, click **Restart**. Log on after the restart.
26. Now create the IDPS rule. At the desktop, click the **Dash home** icon. In the Search box, type **edit**. Click the **Text Editor** icon in the lower window.
27. In the Untitled Document 1 – gedit window, carefully type the following rule:

    alert tcp any any <> any any (content:"password"; msg:"Potential Password Violation"; sid: 11995522;)

    Be sure not to add any line breaks by pressing Enter; Snort rules must be on a single line. If the rule is too long, it will wrap to the next line automatically. When you finish, click the **Save** button. In the Save As window, type **password.rules** in the Name box. In the Places frame, click the name of your user, and then click **Save**. Close the editor.
28. Open a terminal window, type **su root** and press **Enter**. Type your password, and press **Enter**.
29. You should be in your user's home directory, as indicated by the prompt. Type **mv password.rules /etc/snort/rules** and press **Enter**.
30. Create a new log folder. Type **mkdir /etc/snort/log2** and press **Enter**.
31. Log on to Windows 7 with an administrative account. Click **Start**. In the Search programs and files text box, type **\\Team*x*Linux** and press **Enter**. The Windows Security window appears. In the User name text box, type **Team*x*Linux\linux*x*user**; in the Password text box, type **Pa$$word**. Click **OK**. Figure 8-19 shows the configuration of the authentication window.

Used with permission of Microsoft Corporation

Figure 8-19 Authentication from Windows to Linux

32. You should now see the linux*x*user shared folder. Do not do anything more until your partner on Linux completes the next step.
33. On the Linux system, type the following command at the terminal window:

    snort –dev –c /etc/snort/snort.conf –l /etc/snort/log2 –K ascii

    (The –e option between the –d and –v is used to include the Ethernet header, which includes Data Link layer MAC addresses.)

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

34. If you receive errors about rules failing, you can edit /etc/snort/snort.conf and place the # sign at the start of the line that references the failing rule.
35. Once Snort is running on the Linux system, open the linux*x*user shared directory on the Windows 7 system, right-click the white space in the window, click **New**, and click **Text Document**. Name the document **system**. In the body of the new document, enter the word **password**, and then save the document.
36. On the Linux system's terminal window where Snort is running, press **Ctrl+C**.
37. Investigate the /etc/snort/log2 directory. You should find a directory named after the IP address of Windows 7 that contains a file named "TCP:*remote system port – local system port.*" Open this file and examine its content. Open and examine the "alert" file.
38. Log off both systems.

Case Projects

Case Project 8-1: Inline vs. Passive IDPSs

You work for a network solutions provider. Your goal is to become a network engineer and design enterprise networks. You began with the company as a technical support analyst who helped clients troubleshoot hardware and software. Because of your strong communication skills, your boss has been sending you on sales calls, but you feel that this task is moving you away from your goal of network engineering. As a compromise, you and your boss have agreed on a plan in which you will make a series of video training programs to explain complicated networking topics to clients.

Your first video will explain the difference between inline and passive IDPS devices. In preparation for video production, you must submit an outline of your presentation along with explanatory diagrams. Prepare a one- to two-page outline that defines inline and passive IDPSs, the hardware required, the advantages and disadvantages of each, and typical situations in which one is preferred over the other. Include at least two diagrams.

Case Project 8-2: Creating a Snort Rule

You have been investigating an apparent internal attack against your company's Windows Server 2008 file servers. Suspicious packets have been captured during routine audits. You need to configure Snort to log these suspicious files. Your internal network address is 172.20.0.0 with a subnet mask of 255.240.0.0. Your file servers' addresses are 172.20.0.12 and 172.20.0.13. Each of these file servers is running Snort as an HIDPS.

The suspicious packets have the following characteristics:

- They have come from different systems inside your network.
- The packets all include the word *release* between the 1000th and 1100th bytes.
- The packets use TCP as their Transport layer protocol.
- The packets appear to be trying to exploit vulnerabilities in the Windows implementation of SMB over IP.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

You need to write a rule to be included in the rules directory of each server's Snort installation. These two rules must be as specific as possible so that the system logs only the packets that meet the signature of the suspicious network activity. The logged packets should be labeled as "Possible Internal SMB over IP Attack." You must perform research beyond the scope of this chapter to find the needed information and create the rules.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 9

Firewalls

After reading this chapter and completing the exercises, you will be able to:

- Explain what firewalls can and cannot do
- Compare hardware and software firewalls
- Describe common approaches to packet filtering
- Design packet filtering and firewall rule sets

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

When you travel by air, you have to pass through several security checkpoints. Your bags are checked, you go through a scanner, and when you board the plane, your tickets and identity are verified. The purpose of these checks is to protect security and keep unauthorized people off the plane.

In the same way, security devices are placed at a network perimeter to filter packets of digital information. This arrangement of security devices is collectively called a firewall or firewall perimeter. This chapter discusses the concept of firewalls—what they are and what they are not. Then you learn about packet filtering and developing a rule base to carry out your security policy. You also learn about different types of firewall hardware and software, and how to work with proxy servers to make network defense more effective. Finally, you learn how to create a rule base for your firewall.

Overview of Firewalls

A **firewall** is hardware or software that can be configured to block unauthorized access to a network. *Firewall* is often viewed as a catch-all term to represent any device that can block attackers or viruses. Managers commonly ask their network administrators to "get a firewall and put it on the network so that we have better security." This request implies that a firewall is one device that can single-handedly keep attackers away from a network, which is an incorrect assumption. Firewalls cannot protect against employees who send proprietary information outside an organization through a network connection or copy confidential data to a disk. A strong security policy and access controls on sensitive information are needed to protect against such incidents. A firewall also cannot protect connections that do not go through it, such as remote dial-up connections. These connections require a virtual private network (VPN) for protection. The following sections explain what firewalls are and are not so that you have a clearer idea of what they do.

The term *firewall* does not necessarily refer to a single router, computer, VPN gateway, or software program. Any network firewall is a combination of software and hardware components. A firewall can consist of all devices positioned on the network perimeter, so the term *firewall perimeter* might be more descriptive.

The earliest firewalls were packet filters; a single packet-filtering router was placed at the network perimeter to provide protection, although it was rather weak. Today's firewalls are more than just a hardware appliance, and many are strictly software. Some firewall programs are designed for general consumer use, such as ZoneAlarm, which you used in Chapter 1. These programs have relatively simple interfaces that give nontechnical users a few configuration options. For example, Figure 9-1 shows the General Settings tab of the Norton Security Suite Firewall.

With personal firewalls, much of the work of establishing rules and blocking traffic can be done on a case-by-case basis: When the firewall encounters a type of traffic it does not recognize, it prompts you to decide whether the traffic should be blocked or allowed. This process can seem tiresome until the firewall learns what traffic is acceptable, but it is worth the effort.

For example, Figure 9-2 shows the File Insight feature of the Norton Security Suite Firewall. The window reports that the file being downloaded, ipscan22.exe, has a good trust rating and gives the user the option to trust the file. In this case, the user is informed that fewer

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

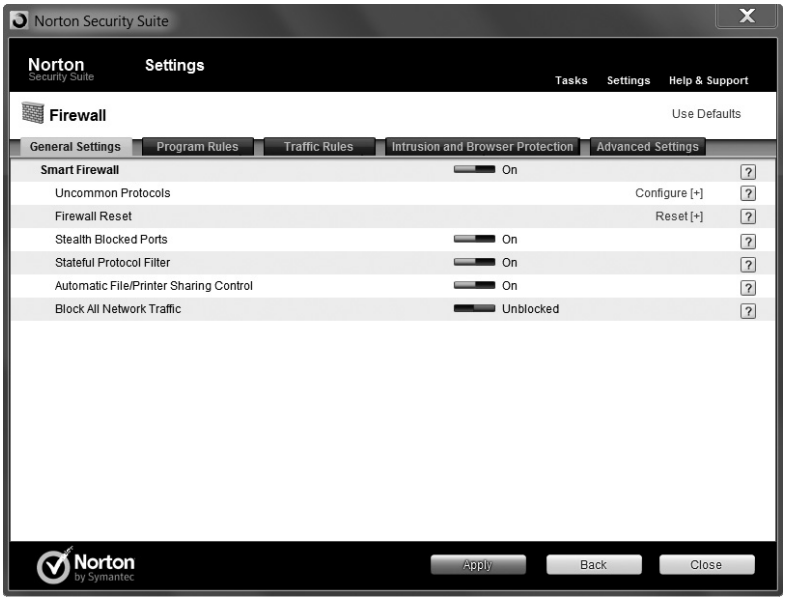

Source: Norton

Figure 9-1 General Settings tab of Norton Security Suite Firewall

than 100 Norton users have downloaded the file, and that it was released only eight days earlier. If the user clicks the Trust Now link, the firewall will remember the setting and apply it automatically in the future.

Source: Norton

Figure 9-2 File Insight window of Norton Security Suite Firewall

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Other firewall programs, such as Check Point NGX, are designed to protect and monitor large-scale networks. They come with a variety of tools for configuring and monitoring network traffic. The range of options in these programs can be quite complex, and many offer integrated VPN, firewall, intrusion detection, and proxy services in one product. Still other **firewall appliances**, such as those in the Cisco PIX line, are self-contained hardware devices with firewall functions you can add to a network.

Whether they are hardware or software, firewalls are effective only if they are configured correctly to block undesirable traffic, such as a SYN flood or an IP spoofing attack. The most sophisticated firewalls are useless if they have poorly configured rules that allow traffic from untrusted sources. In addition, because new threats and modes of attack are devised constantly, firewalls can be configured correctly but still fail to block harmful traffic because they do not yet recognize the traffic as a threat.

Firewalls are not a stand-alone solution, and should not be regarded as the only tool needed to protect a network. No firewall can protect a network from internal threats, such as disgruntled employees or users who try to find ways around security measures. A strong security policy and employee education are essential. Your organization's security policy should also include strict procedures for keeping patches updated and checking for new vulnerabilities. Many worms and viruses have gotten past expensive firewalls that were not patched in a timely manner.

Ideally, a firewall should be combined with antivirus software and an intrusion detection and prevention system (IDPS) for a more comprehensive solution. A strong network security architecture encompasses many components, including an IDPS, firewalls, antivirus software, access control, and auditing. This defense in depth (DiD) approach is the most effective method of securing resources, so keep layered defense strategies in mind. For example, you can integrate the following products with a firewall:

- eSafe Content Security, a product by SafeNet that scans traffic for malware and application-level threats and filters content to improve productivity. It also helps prevent data leaks.
- ManageEngine's Firewall Analyzer, a Web-based monitoring and log-analysis tool that collects logs from proxy servers, firewalls, IDPSs, and VPNs to monitor performance, audit traffic, and detect intrusions

The expense of purchasing and installing firewall software programs can be considerable. The most important consideration is the need for an integrated security system in which a firewall plays an important—but by no means solitary—role.

Comparing Software and Hardware Firewalls

As a network administrator responsible for security, you need to evaluate firewall software and hardware packages and recommend the best choice for your company's needs. Firewalls come in many varieties, but they all handle the core functions of filtering, proxying, and logging. Some packages add caching and address translation. Many vendors now offer firewall packages that also include content filtering, antivirus protection, and intrusion detection. Security products can be free or extremely expensive, but do not let the price rule your decision. Some freeware products have the same functions as more expensive commercial packages, and combinations of freeware and open-source security products can be effective. The following sections

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

describe the software and hardware combinations you are likely to encounter in the workplace. This overview can help you choose a package when the time comes.

Software-Based Firewalls

Most people are familiar with software firewalls, which can be combined with hardware devices to create a secure checkpoint. The downside of software-based firewalls is that they require extensive work to configure the software and secure the operating system (OS) by patching or removing vulnerable services. On the other hand, they are usually less expensive than hardware firewalls, so using them in several locations is affordable. Security professionals often debate whether software firewalls are more secure than hardware firewalls. Many believe the level of security depends more on the administrator's skill in configuring the firewall and the regularity of maintenance and updates than on a specific program or device.

Free Firewall Programs Free firewall programs have some drawbacks. Their logging capabilities are not as robust as those in some commercial products, and configuration can be difficult. In addition, they might not have the ability to monitor traffic passing through the firewall in real time or manage firewall settings for a network from a central location. Nonetheless, they have a place in small-business and home networks because of their convenience and simplicity. The following free firewall programs are popular:

- *Netfilter*—This program comes with the Linux 2.4 and 2.6 kernels and is a solution for stateless and stateful packet filtering, Network Address Translation (NAT), and packet processing. Netfilter logs detailed information about traffic in a well-organized manner that is easy to review.
- *ZoneAlarm*—The free version of this program (which you installed in Chapter 1) is so effective that before configuration is completed, you might lose your Internet connectivity when you first set it up. A correct configuration that allows only the software and IP addresses you want is critical to maintaining Internet connectivity.

Many freeware firewall programs are stripped-down versions of more full-featured products from the same companies. You can try ZoneAlarm Free Firewall, for example, and then decide whether to purchase the commercial version, ZoneAlarm PRO Firewall.

Commercial Firewall Software: Personal Firewalls Personal firewall products are located between the Ethernet adapter driver of the machine where they are installed and the TCP/IP stack so they can inspect traffic between the driver and the stack. Often they are packaged with additional security software. They include programs such as the following:

- *CA (Computer Associates) Internet Security Suite*—This program includes a personal firewall with security levels for grouping programs, protocols, and ports; automated configuration and help for novice users; and features that offer experienced users fine-tuned control. Other features of the suite include cloud-based scanning, browser security, and protection against viruses, spyware, and spam.
- *Norton Internet Security*—This program is designed for home users and includes a personal firewall with controls for both beginners and advanced users. The

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

program includes identity protection along with antivirus, antiphishing, and anti-spyware features.

- *ZoneAlarm Internet Security Suite*—This program includes antivirus, antispam, and spyware protection and the capability to block pop-up ads and cookies.
- *Kaspersky Internet Security*—This highly regarded product includes a personal firewall, identity protection, antivirus protection, cloud-based scanning, rollback of malware changes, antispam protection, and rootkit detection.

Despite the features in the preceding list, these products do not offer extensive firewall protection. Some work with multiple protocols, but most guard only against IP threats. Most assume the responsibility for determining an appropriate level of system security, which makes it easy for a less knowledgeable user to choose a protection level but limits the ability of an experienced user to fine-tune the configuration.

Commercial Firewall Software: Enterprise Firewalls Enterprise firewall programs come with a centralized management option and sometimes have the capability to install multiple instances from a central location. The following programs offer enterprise-level protection:

- *Check Point NGX*—This suite of management and configuration tools offers real-time monitoring, remote administration, and log file analysis.
- *Proventia security products*—This program from IBM Internet Security Systems includes a variety of security tools and packages with an emphasis on centralized management for enterprise applications.

Some firewall products once consisted mostly of packet-filtering tools, but they have added user authentication, NAT, encryption, and centralized management to stay ahead in an increasingly competitive market. As attack profiles move toward Web applications, social media, instant messaging, and streaming media, integrating additional security and management features is becoming the rule rather than the exception.

Hardware Firewalls

One advantage of hardware firewalls is that they do not depend on a conventional OS, such as Windows or Linux, which can leave security vulnerabilities caused by bugs or other flaws. On the other hand, hardware firewalls do run on an OS—some Cisco hardware firewalls run on Cisco's Internetwork Operating System (IOS), for instance—and they can be subject to the same flaws as computer OSs. Hardware firewalls are usually more scalable than software firewalls, and they can handle more data with faster throughput. On the downside, hardware firewalls tend to be more expensive. The following products are available:

- *Cisco ASA series*—Cisco is the world leader in the production of network infrastructure devices, and its ASA (Adaptive Security Appliances) line supports small offices and large enterprise organizations alike. These devices range from the ASA 5505, which supports a maximum throughput of 150 Mbps and 10,000 connections, to the ASA 5585-X, which supports a maximum throughput of 40 Gbps and 10 million connections. ASA devices can provide an IDPS, VPN, content security, and other services.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *Fortinet FortiGate series*—Another multifunction firewall series, the FortiGate-40C supports up to 200 Mbps, while the FortiGate-5000 supports 480 Gbps. Other features include antivirus protection, data loss prevention, and VPN services.
- *Barracuda NG Firewall*—Like the previous products, the Barracuda NG Firewall series includes different models based on an organization's size. For small offices, the Model F10 offers 150 Mbps throughput and 8000 connections. For enterprise datacenters, the Model F900 offers 21 Gbps and 1 million connections. The models include features such as VPN services, traffic flow control, content security, Web and spam filters, NAT, and an IDPS.

Which firewall product is best for your needs? Your choice depends on the number of users you need to protect, the amount of network traffic passing through the firewall, the budget, and the organization's level of concern about security. The budget should not always be the primary consideration; an inexpensive software firewall can become inadequate as a network grows, so buying the strongest one you can afford might be the wisest choice. Table 9-1 lists some advantages and disadvantages to consider when choosing a firewall.

Type of firewall	Advantages	Disadvantages
Software—freeware	Small file size; ease of installation	Only minimal features are offered; lack of technical support
Software—commercial personal firewalls	Simple to install; economical; autoconfiguration features help novice users yet give advanced users more fine-tuned control	Not as full-featured as enterprise products and not as robust as hardware appliances; usually installed on single-computer systems, which reduces security
Software—commercial enterprise firewalls	Usually installed on a dedicated host for maximum security; centralized administration available for large networks; real-time monitoring and other administrative features	Can be difficult to install and configure; tend to be more expensive
Hardware appliances	More scalable than software firewalls; offer faster throughput	Can be expensive and difficult to patch if bugs or security alerts require it

© Cengage Learning 2014

Table 9-1 Firewall advantages and disadvantages

Packet Filtering and Firewall Rule Sets

Chapter 4 introduced you to access control lists (ACLs) and packet filtering on routers. ACLs filter packets by using a rule base to determine whether to allow a packet to pass. In this chapter, you study how to develop a packet-filtering rule base for a firewall. First, you learn about stateless and stateful packet filtering, and then you see how a packet filter's activities depend on other security components and its position in the perimeter's security configuration.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Stateless Packet Filtering

Packet filtering works by screening traffic that arrives on the network perimeter. **Stateless packet filters** determine whether to allow or block packets based on information in protocol headers. Most often, filtering is based on common features of IP protocol headers:

- *IP address*—Each packet-filtering rule specifies a source IP address and a destination IP address.
- *Ports*—Ports give you more control over what you allow and block when creating filtering rules.
- *TCP flags*—Packet filters can allow or block traffic based on the flags in a TCP header.

Table 9-2 shows a simple set of rules for a firewall located at 192.168.120.1 on network 192.168.120.0.

Rule	Source IP	Source port	Destination IP	Destination port	Action
1	Any	Any	192.168.120.0	Above 1023	Allow
2	192.168.120.1	Any	Any	Any	Deny
3	Any	Any	192.168.120.1	Any	Deny
4	192.168.120.0	Any	Any	Any	Allow
5	Any	Any	192.168.120.2	25	Allow
6	Any	Any	192.168.120.3	80	Allow
7	Any	Any	Any	Any	Deny

© Cengage Learning 2014

Table 9-2 Stateless packet-filtering rules

The following list describes the rules in Table 9-2 in more detail:

- *Rule 1*—Many external hosts contacted by a host on the internal network respond by connecting to TCP ports above 1023; this rule enables these connections.
- *Rule 2*—This rule prevents the firewall from connecting to any other hosts, external or internal. The firewall is supposed to monitor traffic, not make connections. An attacker who gets control of the firewall might try to use it to make a connection; this rule blocks that attempt.
- *Rule 3*—This rule provides extra security for the firewall by keeping external users from establishing a connection to it.
- *Rule 4*—This rule enables internal hosts to make connections to computers outside the network.
- *Rule 5*—This rule enables external users to send e-mail into the network. Port 25 is used for SMTP (Simple Mail Transfer Protocol), the protocol for sending e-mail.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *Rule 6*—This rule allows external users to access the network's Web server. Port 80 is used for HTTP (Hypertext Transfer Protocol), the protocol used by Web servers.
- *Rule 7*—This cleanup rule denies any other traffic that has not been explicitly allowed by previous rules.

However, intruders can still get around these defenses. For instance, the problem with Rule 1 is that it also allows faked (spoofed) connections to appear as though they are responding to connections to gain access.

Stateless packet filters do have the advantage of being inexpensive or free. Some are included with routers or open-source OSs. On the other hand, stateless filters can be cumbersome to maintain in a complex network. They are vulnerable to IP spoofing attacks, and they have no form of authentication.

Stateful Packet Filtering

Stateless packet filters are not enough for most organizations because they cannot filter intrusions that occur when someone connects to a computer that has not initiated a connection. Without a connection in place, an attacker can spoof the computer into creating one, resulting in denial of service (DoS) attacks and other attacks. Another potential vulnerability is that a stateless filter handles every packet separately and does not record which packets have already passed through the filter. Packets that were previously forwarded through a connection have no bearing on a stateless filter's decision to forward or drop a packet.

Stateful packet filters, however, keep a record of connections that a host computer has made with other computers by maintaining a file called a **state table**. The packet filter allows incoming packets only from connected external hosts with a record in the state table. For example, a firewall might have Rule 1 from Table 9-2 in its rule base, allowing external connections on ports above 1023. One host on the internal network might have certain connections under way, as shown in the excerpt from a state table in Table 9-3.

Source IP	Source port	Destination IP	Destination port	Connection state
192.168.120.101	1037	209.233.19.22	80	Established
192.168.120.104	1022	165.66.28.22	80	Established
192.168.120.107	1010	65.66.122.101	25	Established
192.168.120.102	1035	213.136.87.88	20	Established
223.56.78.11	1899	192.168.120.101	80	Established
206.121.55.8	3558	192.168.120.101	80	Established
224.209.122.1	1079	192.168.120.105	80	Established

© Cengage Learning 2014

Table 9-3 State table example

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

In the state table, internal hosts on the network with the address range 192.168.120.0 to 192.168.120.255 have established connections with hosts on the external network. Some are contacting HTTP port 80 to connect via the Web, and one has established a connection to an e-mail server listening on port 25. Near the bottom of the table, some external hosts have established connections to hosts on the internal network on port 80.

Suppose that an attacker tries to connect to the network from IP address 201.202.100.1 by sending a packet with the ACK TCP flag. The ACK flag is normally sent at the end of a "handshake" between two networked computers or as part of a routine response to a received packet to indicate that a connection is established. When a stateful firewall receives this packet, it checks the state table for an existing connection between any internal host and the computer at 201.202.100.1. If a connection is not found in the state table, the packet is dropped.

A user-friendly packet filter is built into Windows 7. You access it through the Control Panel's System and Security section, and then click the Windows Firewall link to open the window shown in Figure 9-3. By clicking "Allow a program or feature through Windows Firewall," you can limit the type of program traffic with more precision and specify exceptions, such as ports or programs (see Figure 9-4).

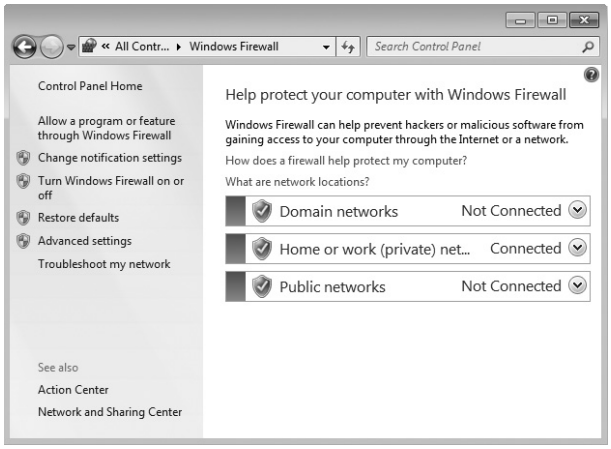

Used with permission of Microsoft Corporation

Figure 9-3 Windows Firewall

By clicking the Advanced settings link in Figure 9-3, you can access Windows Firewall's advanced security settings, as shown in Figure 9-5. Here, you can create very specific rules to allow or deny packets based on protocols, ports, IP addresses, and other items.

Packet Filtering Based on Position

The type of filtering that a firewall, router, or other device performs depends on its position in the firewall perimeter and the other hardware or software with which its activities are coordinated. For instance, a packet filter positioned between the Internet and a host that provides the only protection for that host must be configured carefully; all inbound and outbound traffic should be accounted for in the packet filter's rule base.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

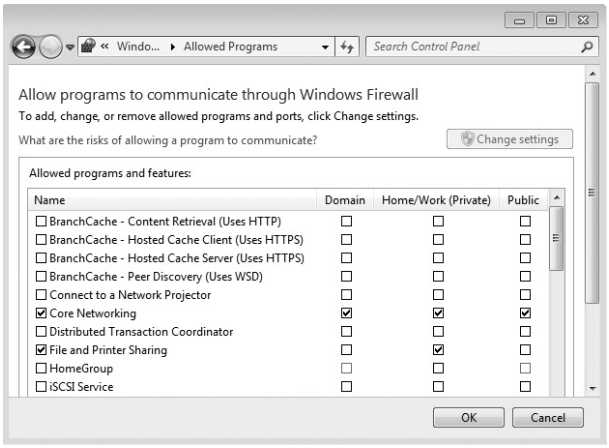

Used with permission of Microsoft Corporation

Figure 9-4 Programs allowed in Windows Firewall

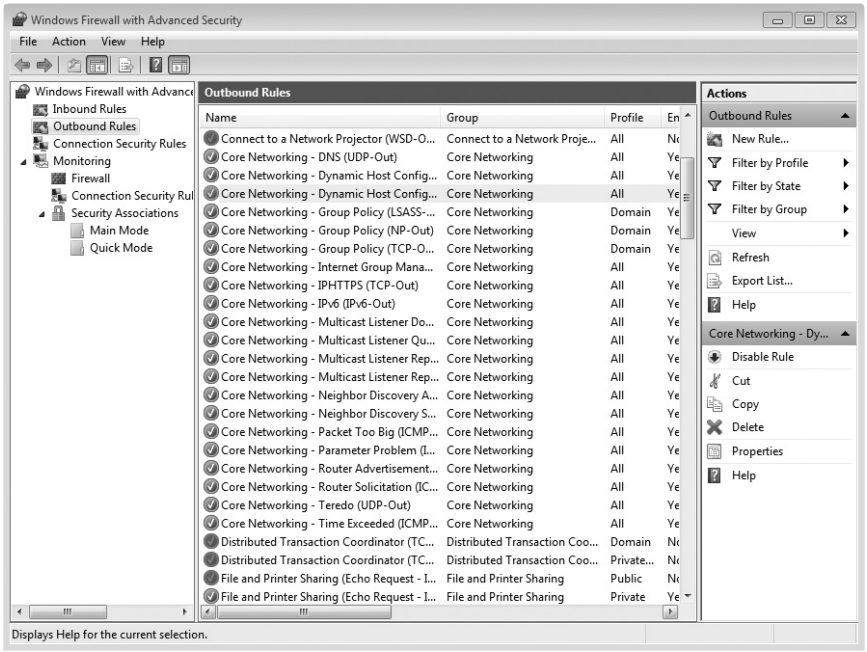

Used with permission of Microsoft Corporation

Figure 9-5 Windows Firewall with Advanced Security window

In contrast, a packet filter placed between a proxy server and the Internet needs to help shield internal users from external hosts. A company concerned about protecting its data warehouses and employee privacy might install a proxy server on the network perimeter. A **proxy server** handles traffic on behalf of computers on the network it protects, rebuilding outbound and inbound requests to hide internal IP address information. A packet filter

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

between a proxy server and the Internet must direct traffic to and from the proxy server. A main goal is to prevent direct connections between the internal network and the Internet (see Figure 9-6).

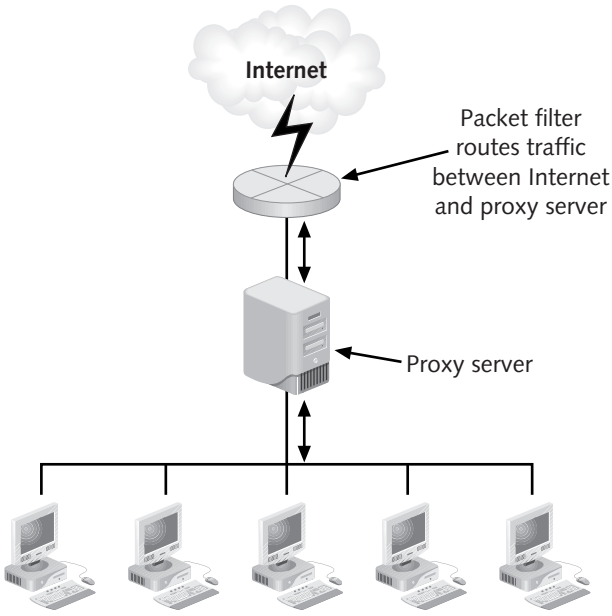

© Cengage Learning 2014

Figure 9-6 A packet filter connecting a proxy server with the Internet

Proxy servers can perform stateful packet filtering on their own. They operate at the OSI model's Application layer, so they can make intelligent decisions about what traffic is allowed to pass. The trade-off is the drop in performance that can result from the demand proxy servers place on the host computer, which is much higher than that of a traditional firewall.

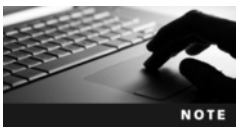

You will learn about proxy servers in more detail in Chapter 10.

Another type of configuration combines packet filtering with a demilitarized zone (DMZ). One common setup is to place packet-filtering devices (routers or firewalls) at either end of the DMZ. The packet filter on the DMZ's external interface needs to allow Internet users to gain access to servers on the DMZ but block access to the internal network. The packet filter on the internal interface performs a similar function for internal users; it enables them to access servers on the DMZ but not connect directly to the Internet. Instead, they connect to the Internet through a proxy server on the DMZ (see Figure 9-7).

The question of how many packet filters to use and where to place them depends on your needs. A simple home network might have adequate protection with a single, well-configured

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

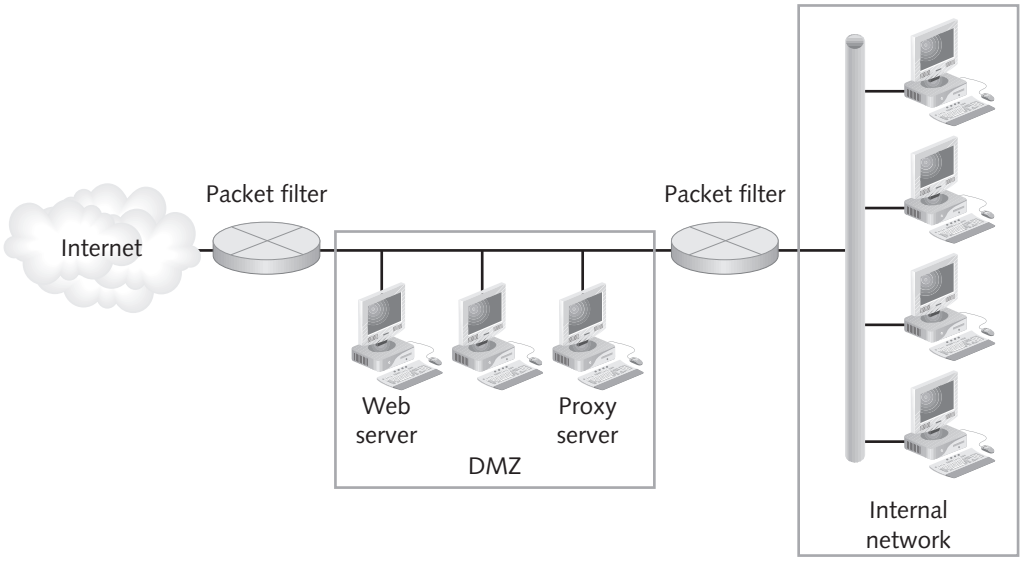

© Cengage Learning 2014

Figure 9-7 A packet filter routing traffic to and from a DMZ

stateful packet filter. A small network that needs to protect proprietary information can use a proxy server-packet filter combination to prevent IP or port information for internal hosts from being revealed to external users. If no IP addresses or ports can be discovered through port scans or other means, no attacks can be launched. For large companies that run public Web servers and have proprietary information to protect, placing packet filters on either side of a DMZ can provide effective multilayered protection for public servers and employees.

Firewall Rule Sets

Most network security professionals agree that even the most full-featured firewall requires a good **rule base**. This set of rules tells the firewall what action to take when a certain kind of traffic attempts to pass through. A simple firewall with well-constructed rules is more effective than a complex product with rules that do not block intrusion attempts.

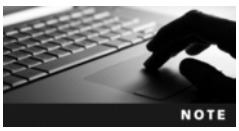

A set of packet-filtering rules is often called an ACL instead of a rule base.

What makes an effective firewall rule base? The following sections describe some points to consider:

- It should be based on the organization's security policy.
- It should include a firewall policy with rules for how applications can access the Internet.
- It should be as simple and short as possible.
- It should restrict access to ports and subnets on the internal network from the Internet.
- It should control Internet services.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Base the Rule Base on Your Security Policy When you are configuring a firewall rule base, you have the responsibility of putting your organization's security policy rules and procedures into practice. Certain elements of filtering packets are especially important to configuring rules and implementing the organization's security policy:

- *Logging and auditing*—Most security policies require methods for detecting intrusions and other security problems, such as viruses.
- *Tracking*—The rule base should include a procedure for notifications so that you can follow response procedures in case of intrusions.
- *Filtering*—A primary objective of a rule base is to filter communications based on complex rules so that the only allowed traffic uses approved protocols, ports, and source and destination IP addresses.
- *NAT*—The rule base should provide for concealing internal names and IP addresses from users outside the network.
- *Quality of Service (QoS)*—QoS rules can be set up to enable the firewall to maintain a baseline level of functionality. This baseline can be specified in the organization's security policy.
- *Desktop security policy*—This policy enables you to specify the level of access that remote users have when they log on to the network.

The following common guidelines should be reflected in the rule base to implement an organization's security policy:

- Employees can have access to the Internet with certain restrictions, such as content filtering or controls on downloads.
- The public can access the company's Web server and e-mail server.
- Only authenticated traffic can access the internal network.
- Employees are not allowed to use instant-messaging or social networking software outside the internal network.
- Traffic from the company's ISP should be allowed.
- External traffic should be blocked if it attempts to connect to a port used by instant-messaging software.
- Only the network administrator should be able to access the internal network directly from the Internet for management purposes.

Create a Firewall Policy for Application Traffic A **firewall policy** is an addition to the security policy that describes how firewalls should handle application traffic, such as that from Web and e-mail applications. A risk assessment should provide a list of these applications and their threats and vulnerabilities, describe the impact if confidential data is compromised, and outline countermeasures for mitigating risks. Before developing a firewall policy, you must understand these factors because risk assessment results dictate whether firewalls block or allow your network's application traffic.

The firewall policy should explain how the firewall is set up, managed, and updated. Although each organization's needs are different, you follow a general set of steps to create a firewall policy. First, you identify which network applications are needed and

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

the vulnerabilities associated with each application. Next, you need to conduct a cost-benefit analysis to determine the most cost-effective and efficient method for securing application traffic. Remember that some traffic is more sensitive than other forms, and you must balance security, user requirements, and cost. Also, keep in mind that many networks have multiple firewalls at several locations, so you need to develop a traffic matrix for each location. Table 9-4 shows an example.

Application or service	Internal host type	Location	Host security policy	Firewall internal security policy	Firewall external security policy
FTP	Windows	Any	Client only; antivirus	Allow	Deny
FTP	UNIX	Any	Secure Shell (SSH); user ID/password; no anonymous traffic	Allow	Application proxy with user authentication
Telnet	Windows	Any	Client only	Allow	Application proxy with user authentication
Telnet	UNIX	Any	SSH	Allow	Application proxy with user authentication
SMB over IP	Windows	Any	Limit access to shares	Allow local domain only; deny all others	Deny

© Cengage Learning 2014

Table 9-4 Application traffic matrix

When the traffic matrix is finished, you can develop the firewall rule base. Firewalls enable you to control access to your computer or network by controlling access to particular applications. When computers listen for incoming connections on ports used by instant-messaging programs and other applications, attackers can scan for these open ports. Firewall software allows you to block applications from accessing the Internet so that these ports are not left open. When one of these applications needs to be used, the user can be prompted to specify whether access to the Internet is allowed. Typically, you have three options:

- *Allow*—The application can access the Internet at any time. You will probably want to establish this rule for Web browsers, e-mail applications, and other software used frequently.
- *Block*—The application is blocked from accessing the Internet.
- *Ask or prompt*—The user is prompted when the application attempts to access the Internet.

Keep the Rule Base Simple Keep the list of rules in your rule base as short as possible. The more complex the rule base is, the greater the chance of introducing errors. Some professionals suggest that a rule base contain no more than 30 rules, and certainly no more than 50. Also, the shorter the rule base is, the faster your firewall can process requests because it has fewer rules to review.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Typically, a firewall processes rules in a particular order. The rules are usually numbered sequentially, and the list is usually displayed in the form of a grid. The first cell in the grid is the rule number. Subsequent cells describe attributes the firewall should test for, such as an IP address or a protocol, and an action to be taken. Because rules are processed in order, the most important ones—the ones that should be processed first—are placed at the top of the list. (All firewalls have a way to reorder rules in the rule base.) It is a good idea to make the last rule a **cleanup rule**, which handles any other packets that have not been covered in preceding rules. On a Cisco router, this cleanup rule is the implicit "deny all" statement at the end of an ACL. By default, this rule means that all other traffic not addressed in the rule set will be denied.

When a packet hits a firewall, rules are tested one after another. As soon as a match is found, the prescribed action is taken. Usually, the two possible actions are to allow or deny the packet. The Allow action permits the packet to proceed through the firewall to its destination on the Internet or in the internal network; the Deny action drops the packet. Usually, no notification is sent to the sender for a Deny action because it might give attackers clues about the network's characteristics; the packet is simply dropped. A third option, Reject, typically is not used because it notifies the sender and could provide clues about the network.

Restrict Subnets, Ports, and Protocols Packet filters can usually give you more control over what traffic you block or allow. The more flexibility you have, the better you can control specific types of traffic. The following sections provide some examples of rules you can create to control traffic yet allow necessary connectivity.

Filtering by IP Addresses One way to identify traffic is by IP address range. Some traffic comes from networks you trust, such as your own network or your ISP's servers, so it should not be blocked by the firewall. However, most firewalls start from the sensible premise that all traffic should be blocked by default, and you need to identify "trusted" networks whose communications should be allowed.

For instance, if your network subnet contains IP addresses in the range of 172.20.0.1 to 172.20.2.254, you would specify that the firewall regard them as "trusted," as shown in one of ZoneAlarm's configuration windows (see Figure 9-8). Remember that 172.20.0.0 is the subnet address and 172.20.2.255 is the broadcast address for this range; these addresses cannot be used as valid IP addresses for hosts.

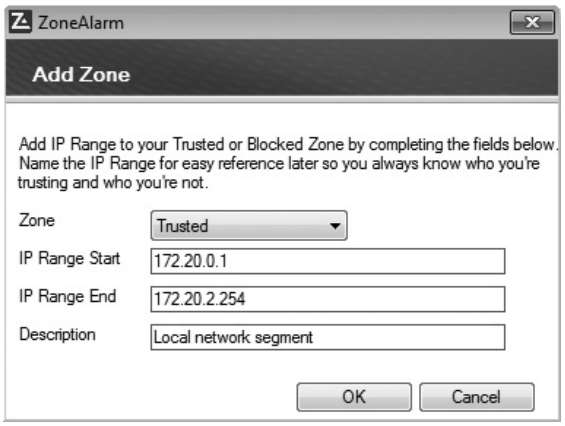

Source: ZoneAlarm

Figure 9-8 Identifying trusted subnets or IP addresses

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Filtering by Ports TCP/IP and UDP transmit information by breaking it into sections called segments (TCP) or datagrams (UDP and IP). While packets can vary in length, it is important to remember that a system using TCP or UDP sends packets of the same length so that they can be sent and received more easily. A file might be divided into several packets, but as many of the packets as possible are the same length.

Filtering by TCP or UDP port numbers is commonly called port filtering or protocol filtering. Using TCP or UDP port numbers can help you filter a wide variety of information, including SMTP and POP e-mail, SMB (Server Message Block) sessions, Domain Name System (DNS) requests, and Network News Transfer Protocol (NNTP) newsgroup sessions. You can filter out everything but TCP port 80 for Web traffic, TCP port 25 for e-mail traffic, or TCP port 21 for FTP transmissions. Figure 9-9 shows the port-filtering process.

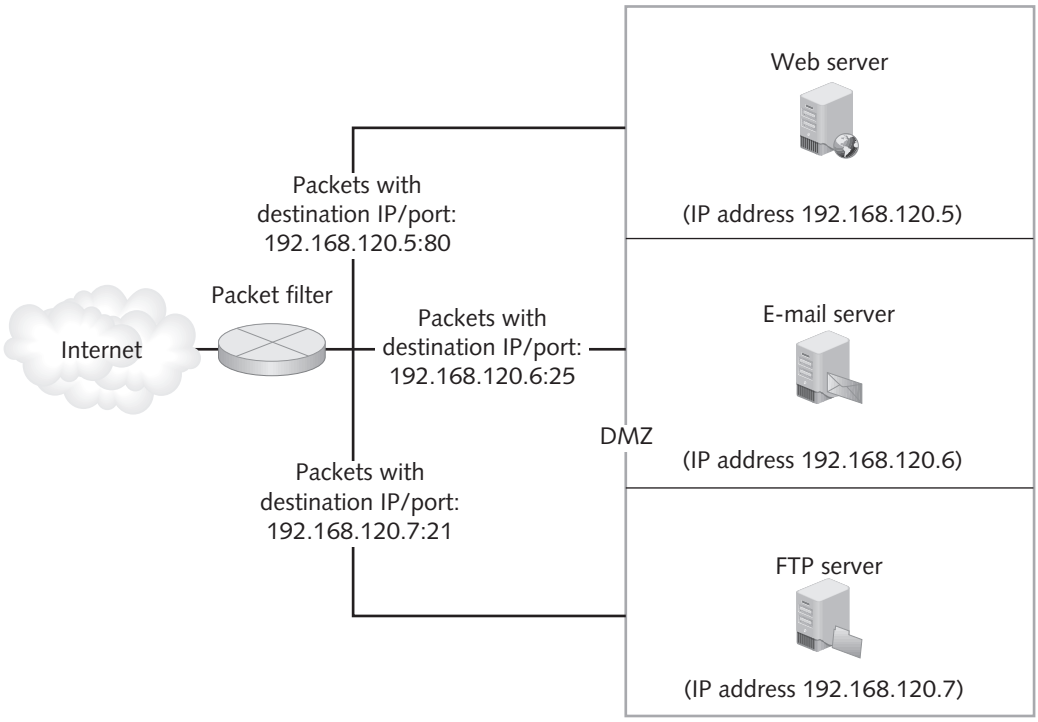

© Cengage Learning 2014

Figure 9-9 Port numbers direct packets to the client or server that needs them

Figure 9-9 indicates that ports are filtered by a router and directed to the correct combination of computer and software by their IP address and port number. The combination of an IP address and a port number is called a **socket**. It is usually shown with a colon separating the IP address from the port number—for example, 192.168.120.5:80. The port numbers shown in the figure are technically correct; port 80 is reserved for HTTP Web pages, port 25 for SMTP e-mail, and so on. However, in practice, two computers exchange information using two different port numbers: the port from which data is sent (the source port) and the port where the response data is received (the destination port). Network communication

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

between two computers commonly uses different source and destination ports; in fact, the ports are rarely the same. However, this approach makes filtering by port number complicated. Another complicating factor is that the destination port is determined dynamically on a per-connection basis, making it impossible to predict.

There is no easy solution when filtering services based on port number, especially when the ports are above 1023 in the registered or dynamic port ranges that IANA assigns. You must know what services you are running on your network and what ports they use; you block the ports you do not need and allow the ones you need.

For instance, on a Windows-based network, computers probably use the NetBIOS name service to find one another on NetBIOS ports 137, 138, and 139. NetBIOS traffic should originate inside your internal network and stay inside the firewall; any traffic from the Internet that attempts to use these ports should be dropped. You should also block inbound traffic that attempts to use ports assigned to other Windows networking services. Table 9-5 lists some examples of filtering common Windows services and ports. These filters would apply to traffic trying to enter the organization's network from the Internet.

Rule	Protocol	Source IP	Source port	Destination IP	Destination port	Action	Service
1	TCP	Any	Any	Any	135	Deny	NetBIOS RPC
2	UDP	Any	Any	Any	135	Deny	NetBIOS RPC
3	TCP	Any	Any	Any	137	Deny	NetBIOS Name Service
4	UDP	Any	Any	Any	137	Deny	NetBIOS Name Resolution
5	UDP	Any	Any	Any	138	Deny	NetBIOS Datagram Service
6	TCP	Any	Any	Any	139	Deny	NetBIOS Session Service
7	UDP	Any	Any	Any	161	Deny	SNMP
8	TCP	Any	Any	Any	445	Deny	SMB/file sharing
9	UDP	Any	Any	Any	1755	Deny	Windows Media
10	TCP	Any	Any	Any	3389	Deny	Remote Desktop Protocol
11	UDP	Any	Any	Any	3389	Deny	Remote Desktop Protocol

© Cengage Learning 2014

Table 9-5 Filtering Windows services and ports

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

In Table 9-5, RPC stands for Remote Procedure Call. RPC is a critical component of Windows communications, so be careful when filtering it. Be sure to check for dependencies, as you should with any service or protocol you are filtering. SNMP stands for Simple Network Management Protocol.

UNIX has its own set of services that should be blocked when they are inbound from the Internet. Table 9-6 lists some examples.

Rule	Protocol	Source IP	Source port	Destination IP	Destination port	Action	Service
1	TCP	Any	Any	Any	17	Deny	QOTD
2	UDP	Any	Any	Any	17	Deny	QOTD
3	TCP	Any	Any	Any	111	Deny	Portmapper
4	UDP	Any	Any	Any	111	Deny	Portmapper
5	TCP	Any	Any	Any	513	Deny	Remote Login
6	TCP	Any	Any	Any	514	Deny	Syslog
7	UDP	Any	Any	Any	514	Deny	Syslog
8	TCP	Any	Any	Any	635	Deny	mountd (NFS service)

© Cengage Learning 2014

Table 9-6 **Filtering UNIX services and ports**

In Table 9-6, QOTD stands for Quote of the Day. Portmapper is similar to RPC. Syslog is the system logging service, and mountd is used for mounting drives in the NFS (Network File System), which is similar to SMB.

The problem with specifying a port number for filtering is that some applications, particularly instant-messaging programs, do not use a fixed port number when they reply to a client; they assign themselves dynamic port numbers, which are used only for the length of a particular communication. You have to filter these communications by IP address instead.

Filtering by Service With some firewalls, you can filter by naming the service you want to use instead of specifying a port number. For example, instead of filtering port 23, you can simply specify the Telnet service. Packets can also be filtered based on the ID field in the IP header. This method works in much the same way for firewalls as it does for IDPSs. Firewalls can also filter by the six TCP control flags you learned about in Chapter 2:

- URG (urgent)
- ACK (acknowledgement)

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- PSH (push function)
- RST (reset the connection)
- SYN (synchronize sequence numbers)
- FIN (no more data from sender)

You have also learned about IP options, the set of optional flags in an IPv4 header. Options add information to a packet, as in these examples:

- *Security*—This option enables hosts to send security parameters, handling restrictions, and other information.
- *Loose source and record routing*—This option enables the source computer to specify the routers used in forwarding the packet to its destination.
- *Strict source and record routing*—This option is the same as loose source and record routing, except that the host computer must send the packet directly to the next address in the source route.
- *Internet timestamp*—This option provides a timestamp indicating when the packet was sent.

IPv6 presents other issues for firewalls. For example, IPv6 contains a large number of error messages, informational messages, and responses, such as Neighbor Discovery, Router Discovery, and Stateless Autoconfiguration. This information can be useful to attackers and must be prevented from leaving the internal network, but it must pass through internal firewalls.

ICMPv6, like its predecessor ICMP (Internet Control Message Protocol), is very useful for network diagnostics, but you must draw a fine line in security and firewall policies to be sure that risk is minimized when this type of traffic is allowed to leave the internal network. Less strict policies can be established for traffic that is authenticated, such as an IPsec AH header or ESP with authentication.

Packet-filtering rules should be tailored to meet a network's needs. However, every rule base should follow a few general practices:

- A firewall or packet filter that follows a "deny all" security policy should begin by allowing services selectively as needed and end by blocking all other traffic.
- The rule base should keep everyone except network administrators from connecting to the firewall. Anyone who accesses the firewall could discover internal IP addresses and gain access to the internal network.
- The rule base should block direct access from the Internet to any computers behind the firewall or packet filter. All inbound traffic, in other words, should be filtered first.
- The rule base should permit access to public servers in the DMZ and enable users to access the Internet.

A rule base, then, is a mixture of rules that selectively allow or deny access. Table 9-7 lists a typical set of rules that assume the firewall is at 192.168.120.1, the e-mail server is at 192.168.120.2, the Web server is at 192.168.120.3, and the DNS server is at 192.168.120.4. The internal network is represented by 192.168.120.0.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Rule	Source IP	Source port	Destination IP	Destination port	Action	What it does
1	192.168.120.1	Any	Any	Any	Deny	Prevents the firewall from making any connections
2	Any	Any	192.168.120.1	Any	Deny	Prevents anyone from connecting to the firewall
3	192.168.120.0	Any	Any	Any	Allow	Allows internal users to access external computers
4	192.168.120.0	Any	192.168.120.4	53	Allow	Enables internal users to connect to the DNS server on port 53
5	Any	Any	192.168.120.2	25	Allow	Allows external and internal users to access the e-mail server via SMTP port 25
6	192.168.120.0	Any	192.168.120.2	110	Allow	Enables internal users to connect to the e-mail server using POP3 port 110
7	Any	Any	192.168.120.3	80	Allow	Enables both external and internal users to connect to the Web server on HTTP port 80
8	Any	Any	Any	Any	Deny	Blocks all traffic not covered by previous rules

© Cengage Learning 2014

Table 9-7 A typical packet-filtering rule base

Control Internet Services The goal of packet-filtering rules is to account for all possible ports a type of communication might use or for all variations of a protocol (for instance, passive and active FTP or standard and secure HTTP). Most businesses use the Internet heavily, and the challenge for security professionals is balancing the need for Internet connectivity with the need to secure network resources from the risks of Internet access. The following sections discuss how to control access to commonly used Web resources, such as the Internet, sending and receiving e-mail, and using FTP as a secure means of centralized data storage.

Web Services Rules A common priority for employees in a protected network is the ability to use the Internet and exchange e-mails. Rules for Web access need to cover standard HTTP traffic on TCP port 80 and Secure HTTP (S-HTTP) traffic on TCP port 443. These rules, shown in Table 9-8, assume that the local network has an IP address range of 207.177.178.0/24.

DNS Resolution To connect to Web sites, users need to be able to resolve the fully qualified domain names (FQDNs) they enter, such as www.cengage.com, to their corresponding

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Rule	Protocol	Transport protocol	Source IP	Source port	Destination IP	Destination port	Action
1	HTTP outbound	TCP	207.177.178.0/24	Any	Any	80	Allow
2	S-HTTP outbound	TCP	208.177.178.0/24	Any	Any	443	Allow

© Cengage Learning 2014

Table 9-8 Outbound Web access

IP addresses by using DNS. Internal users may connect to external hosts via a DNS server in the DMZ. DNS uses UDP port 53 for name resolution attempts, and uses TCP port 53 if an initial UDP query fails and for zone transfers. In addition, you need to set up rules that enable external clients to access computers in your network by using the same TCP and UDP ports, as shown in Table 9-9. The rules in Table 9-9 assume that the DNS server is located at IP address 208.177.178.31.

Rule	Protocol	Transport protocol	Source IP	Source port	Destination IP	Destination port	Action
3	DNS outbound	TCP	208.177.178.31	Any	Any	53	Allow
4	DNS outbound	UDP	208.177.178.31	Any	Any	53	Allow
5	DNS inbound	TCP	Any	Any	208.177.178.31	53	Allow
6	DNS inbound	UDP	Any	Any	208.177.178.31	53	Allow

© Cengage Learning 2014

Table 9-9 DNS resolution rules

E-Mail Configuration To set up e-mail configuration rules, assess whether your organization needs to accept incoming e-mail, whether internal users can access Web-based e-mail (such as Hotmail), and what e-mail clients your organization supports. This information helps ensure security without blocking e-mail access. Setting up firewall rules that permit filtering e-mails is not simple, however. One reason is the variety of e-mail protocols you can use:

- POP3 and IMAP4 for inbound mail transport
- SMTP for outbound mail transport
- Lightweight Directory Access Protocol (LDAP) for looking up e-mail addresses
- HTTP for Web-based e-mail service

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

To keep things simple, consider a configuration that uses only POP3 and SMTP for inbound and outbound e-mail, respectively. However, you should use Secure Sockets Layer (SSL) encryption for additional security; SSL is a form of encryption between Web server and client. Table 9-10 outlines rules for e-mail access, assuming that the SMTP mail server is located at 208.177.178.29. If an *S* follows a protocol, such as POP3/S, the protocol has SSL encryption.

Rule	Protocol	Transport protocol	Source IP	Source port	Destination IP	Destination port	Action
7	POP3 outbound	TCP	208.177.178.0/24	Any	Any	110	Allow
8	POP3/S outbound	TCP	208.177.178.0/24	Any	Any	995	Allow
9	POP inbound	TCP	Any	Any	208.177.178.0/24	110	Allow
10	POP3/S inbound	TCP	Any	Any	208.177.178.0/24	995	Allow
11	SMTP outbound	TCP	208.177.178.29	Any	Any	25	Allow
12	SMTP/S outbound	TCP	208.177.178.29	Any	Any	465	Allow
13	SMTP inbound	TCP	Any	Any	208.177.178.29	25	Allow
14	SMTP/S inbound	TCP	Any	Any	208.177.178.29	465	Allow

© Cengage Learning 2014

Table 9-10 E-mail rules

FTP Transactions Two kinds of FTP transactions can take place on networks: active FTP or passive FTP. The rules you set up for FTP need to support two separate connections: TCP 21, the FTP control port, and TCP 20, the FTP data port. If some clients in your network support active FTP, you cannot specify a port because the client can establish a connection with the FTP server at any port above 1023. Instead, you specify the IP address of your FTP server (in this example, 208.177.178.25), as shown in Table 9-11.

ICMP Message Types ICMPv4 and v6 function like housekeeping protocols for TCP/IP, helping networks cope with communication problems. The downside is that because ICMPv4 packets have no authentication method to verify the packet recipient, intruders can attempt man-in-the-middle attacks in which they impersonate the intended recipient. They can also send ICMP Redirect packets to direct traffic to a computer they control outside the protected network. A firewall/packet filter must be able to determine whether an ICMP packet should be allowed to pass, based on its message type.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Rule	Protocol	Transport protocol	Source IP	Source port	Destination IP	Destination port	Action
15	FTP control inbound	TCP	Any	Any	208.177.178.0/25	21	Allow
16	FTP data inbound	TCP	208.177.178.0/25	20	Any	Any	Allow
17	FTP PASV (passive)	TCP	Any	Any	208.177.178.0/25	Any	Allow
18	FTP control outbound	TCP	208.177.178.0/25	Any	Any	21	Allow
19	FTP data outbound	TCP	Any	20	208.177.178.25	Any	Allow

© Cengage Learning 2014

Table 9-11 FTP rules

Table 9-12 lists some rules you can use to block common ICMPv4 message types. The INT and EXT (internal and external) columns are left empty for Rules 20 and 26 so that those rules can apply to both inbound and outbound traffic. Rule 26 is a cleanup rule that drops all ICMPv4 packets not filtered by previous rules.

Filtering ICMPv6 packets is significantly different from filtering ICMPv4 packets. For one thing, the address space in IPv6 is so large that port scanning is much more difficult for

Rule	INT	EXT	Protocol	ICMP type	Source IP	Source port	Destination IP	Destination port	Action
20			ICMP	Source Quench	Any	Any	Any	Any	Allow
21	X		ICMP	Echo Request	Any	Any	Any	Any	Deny
22		X	ICMP	Echo Reply	Any	Any	Any	Any	Deny
23	X		ICMP	Destination Unreachable	Any	Any	Any	Any	Allow
24		X	ICMP	Redirect	Any	Any	Any	Any	Deny
25		X	ICMP	Destination Unreachable	Any	Any	Any	Any	Deny
26			ICMP		Any	Any	Any	Any	Deny

© Cengage Learning 2014

Table 9-12 Filtering ICMP message types

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

attackers. Also, until IPv6 is universally adopted, transmission methods must be available that support both IPv4 and IPv6. One common method is Teredo tunneling, which encapsulates IPv6 datagrams within IPv4 packets. To support Teredo, many connectivity-checking messages must be allowed through the firewall. You should allow the following ICMPv6 message types to pass through the firewalls within an organization but never outside the organization:

- Destination Unreachable (Type 1)—All codes
- Packet Too Big (Type 2)
- Time Exceeded (Type 3)—Code 0 only
- Parameter Problem (Type 4)—Codes 1 and 2 only
- Echo Request (Type 128)
- Echo Response (Type 129)
- Time Exceeded (Type 3)—Code 1
- Parameter Problem (Type 4)—Code 0

Chapter Summary

- A firewall is hardware or software configured to block unauthorized access to a network. A firewall can also be a combination of software and hardware components. The term *firewall* can refer to all devices positioned on the network perimeter, regardless of whether they are hardware or software based.
- Firewalls are not a stand-alone solution. Strong network security encompasses many components, including IDPSs, firewalls, antivirus software, access control, and auditing.
- Software firewalls come in many varieties, including freeware, shareware, and commercial enterprise applications. Hardware firewall appliances are more expensive, but they can handle more traffic.
- Stateless firewalls filter traffic based on basic parameters such as protocol or IP address, but they are much less secure than stateful firewalls that maintain state tables. State tables are records of connections that enable the firewall to make filtering decisions based on whether a trusted computer initiated a session or whether an unknown host outside the company network is trying to establish a connection.
- Firewalls are effective only if they are configured correctly to block undesirable traffic and allow necessary traffic.
- An effective firewall rule base should be based on the organization's security policy, provide rules for how applications can access the Internet, and be as simple and short as possible. A rule base should also restrict access to ports and subnets on the internal network from the Internet, and it should control Internet services.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Key Terms

cleanup rule A packet-filtering rule that comes last in a rule base and covers any packets that have not been covered by preceding rules.

firewall Hardware or software configured to block unauthorized access to a network.

firewall appliances Hardware devices with firewall functionality.

firewall policy An addition to a security policy that describes how firewalls should handle application traffic, such as Web or e-mail applications.

proxy server Software that forwards network packets and caches Web pages to speed up network performance.

rule base The collection of rules that filter traffic at an interface of a firewall.

socket The end point of a computer-to-computer connection defined by an IP address and port address.

state table A file maintained by stateful packet filters that contains a record of all current connections.

stateful packet filters Filters that are similar to stateless packet filters, except that they also determine whether to allow or block packets based on information about current connections.

stateless packet filters Simple filters that determine whether to allow or block packets based on information in protocol headers.

Review Questions

1. A firewall can do which of the following? (Choose all that apply.)
   a. Screen traffic for viruses.
   b. Determine what user is sending transmissions.
   c. Filter traffic based on rules.
   d. Provide a layer of protection for the network.
2. A firewall is an effective stand-alone security solution. True or False?
3. Stateless packet filters allow or block packets based on which of the following?
   a. status of the connection
   b. information in protocol headers
   c. state table
   d. packets that have been handled previously
4. Which of the following is an advantage of using a software firewall rather than a hardware firewall?
   a. throughput
   b. reliability
   c. cost
   d. availability

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

5. Which of the following is an advantage of using a hardware firewall rather than a software firewall? (Choose all that apply.)
   a. scalability
   b. cost
   c. ease of maintenance
   d. increased throughput

6. Almost every type of firewall depends on what configurable feature for its effectiveness?
   a. network connection
   b. state table
   c. rule base
   d. management console

7. Where should you place the most important rules in a rule base?
   a. in the connection log file
   b. at the bottom of the rule base
   c. in the state table
   d. at the top of the rule base

8. Which of the following is a guideline for developing a firewall rule base? (Choose all that apply.)
   a. The rule base should restrict all Internet access.
   b. The rule base should restrict access to ports and subnets on the internal network from the Internet.
   c. The rule base should be as detailed as possible.
   d. The rule base should not interfere with application traffic.

9. A firewall policy does which of the following? (Choose all that apply.)
   a. describes how employees can use the firewall
   b. identifies and mitigates risks
   c. explains how the firewall is set up, managed, and updated
   d. specifies how the firewall should handle application traffic

10. A rule base should end with a(n) _____ rule.
    a. reject
    b. allow
    c. cleanup
    d. block

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

11. When you request a Web page, which port does the Web server use to send you the page?
    a. 80
    b. 443
    c. one higher than 1023
    d. one lower than 1023
12. Stateless packet filters are more secure than stateful packet filters because they do not contain a state table that can be exploited by an attacker. True or False?
13. A socket is a combination of a(n) _____ and a(n) _____.
    a. NetBIOS name, port number
    b. port number, MAC address
    c. MAC address, IP address
    d. IP address, port number
14. The Windows RPC service works like the UNIX _____ service.
    a. mountd
    b. Portmapper
    c. QOTD
    d. INFS
15. Which port is used for name/address resolution?
    a. 20
    b. 53
    c. 80
    d. 110

Hands-On Projects

Hands-On Project 9-1: Exploring the Advanced Settings of Windows Firewall

Time Required: 15 minutes

Objective: Examine the advanced settings of Windows Firewall.

Description: The Windows Firewall in Windows 7 and Windows Server 2008 R2 has three profiles. From most restrictive to least restrictive, they are Public, Private, and Domain. Only one profile can be active at a time. The Windows Firewall has two interfaces: The interface in the Control Panel is appropriate for inexperienced users, and the advanced settings are appropriate for experienced users and technical staff.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

1. If you completed the Chapter 1 hands-on projects, you installed ZoneAlarm on your Windows 7 system. Log on to the Windows 7 system, click **Start**, click **Control Panel**, click **Programs and Features**, and uninstall ZoneAlarm.
2. Click **Start**, click **Control Panel**, click the **View by** list box in the upper-right corner, and click **Small icons**. Click **Windows Firewall**.
3. Log on to the Windows Server 2008 system, and repeat Step 2 to open **Windows Firewall**.
4. On both computers, the Domain networks profile should be listed as connected. The section should be expanded and show that the Windows Firewall is on and that the Active domain network is your domain. Your results should be similar to those in Figure 9-10.

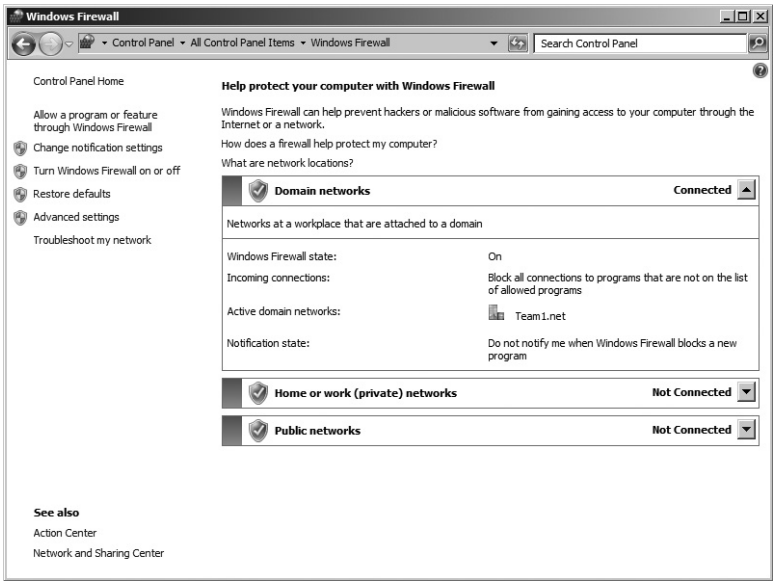

Used with permission of Microsoft Corporation

Figure 9-10 Windows Firewall on a domain computer

5. On both computers, click **Advanced settings**. As usual, when attempting administrative tasks in Windows 7, you need to provide domain administrator credentials when prompted.
6. The Overview section of the Windows Firewall with Advanced Security window should show that the Domain profile is active. The Private and Public profiles show that Windows Firewall is on, but these profiles are not active. In the middle frame, scroll down and click **Windows Firewall Properties**. Here, the first three tabs allow you to customize settings for each of the three profiles. On the Domain Profile tab, click the **Customize** button in the Settings section. In the Firewall settings section, select **Yes** from the Display a notification list box, and click **OK**. In the Logging section, click **Customize**. Note the location of the firewall log. Click **Start**, click **Computer**, and find the folder on the C: drive that holds the log files. The folder should be empty.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

7. Return to the Customize Logging Settings for the Domain Profile window, set both the Log dropped packets and Log successful connections list boxes to **Yes**, and click **OK**. Click **OK** again in the next window.
8. In the left frame of the Windows Firewall with Advanced Security window, click **Inbound Rules**. Compare the differences between the rules listed in Windows Server 2008 and Windows 7. Double-click an inbound rule in the middle pane that is labeled with a green checkmark; notice that the Enabled box is checked. Click **Cancel**, and then double-click an inbound rule labeled with a gray icon; note that the Enabled box is not checked. Click **Cancel**.
9. Leave your systems running for the next project.

Hands-On Project 9-2: Enforcing IPsec Policies

Time Required: 40 minutes

Objective: Increase security for your domain by enforcing encryption and authentication.

Description: While IPsec is commonly used in VPN connections, as you learned in Chapter 5, it can also be used to authenticate and encrypt transmissions between computers on a LAN. A common example pertains to the critical security issue of data being transferred from a publicly accessible Web server on the DMZ to a database server on the trusted internal network that is not publicly accessible. This data often contains highly sensitive information such as customers' billing information; transmitting such information without encryption would be negligent for e-commerce companies. In this project, you require IPsec communications between computers in your domain.

1. Log on to the Windows 7 system as the Team*x* domain user administrator, where *x* represents the number assigned by your instructor.
2. Click **Start**, click **Control Panel**, click **Network and Internet**, and click **Network and Sharing Center**. Click the **Team*x*.net** icon to display the network computers. You should see both TEAM*x*CLIENT and TEAM*x*SRV. Double-click **TEAM*x*SRV**; you should see the default shares (see Figure 9-11).

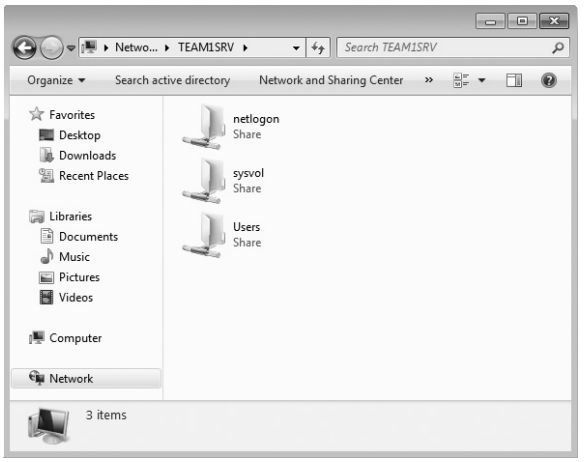

Used with permission of Microsoft Corporation

Figure 9-11 TEAM*x*SRV's default shares

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

3. Double-click the **Users** share, and then drill down through the **Administrator** and **Desktop** folders. In the Desktop folder, right-click the white space in the right frame, click **New**, and select **Text Document**. Type **Private** as the name of the document, and press **Enter**. Close all windows.
4. On the Windows Server 2008 system, close all windows and verify that the Private file has been created on the Administrator's desktop. Next, you will create a firewall rule that restricts any unsecured communications from entering or leaving the server.
5. On Windows Server 2008, open Windows Firewall with Advanced Security. In the left frame, click **Connection Security Rules**. In the right frame, click **New Rule**. In the Rule Type window, verify that **Isolation** is selected, and click **Next**. In the Requirements window, click the **Require authentication for inbound and outbound connections** option button, and click **Next**. In the Authentication Method window, click the **Computer (Kerberos V5)** option button, and click **Next**. In the Profile window, uncheck the **Private** and **Public** boxes, leave the Domain box checked, and click **Next**. In the Name window, type **Domain Authenticated Connections Policy** in the Name text box, and click **Finish**. In the Connection Security Rules frame, verify that the new rule is enabled.
6. From Windows 7, try to complete Steps 2 and 3 again. This time, the server policy prevents the connection. You should receive an error similar to that shown in Figure 9-12 when you try to access TEAM*x*SRV.

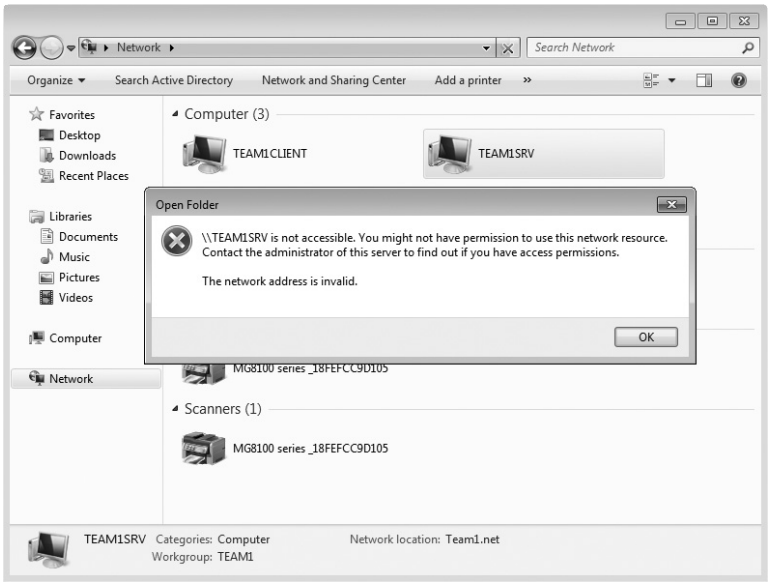

Used with permission of Microsoft Corporation

Figure 9-12 The authentication policy prevents unsecured connections

7. Click **OK** to close the error message, and then close the Network window. Now you will create an IPsec policy that allows Windows 7 to communicate with the server. On the Windows 7 system, click **Start**, click **Control Panel**, click **Administrative**

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Tools, and double-click **Local Security Policy**. Right-click **IP Security Policies on Local Computer**, and click **Create IP Security Policy**.

8. In the welcome window of the IP Security Policy Wizard, click **Next**. In the IP Security Policy Name window, type **Domain Connectivity Policy** in the Name text box, and click **Next**. In the Requests for Secure Communication window, click **Next**. In the Completing the IP Security Policy Wizard window, verify that the Edit properties box is checked, and click **Finish**.
9. Now you need to configure the policy to determine the types of communications to which it applies and what action to take when such communications are detected. In the Domain Connectivity Policy Properties window, click **Add**. In the welcome window of the Security Rule Wizard, click **Next**. In the Tunnel Endpoint window, verify that **This rule does not specify a tunnel** is selected, and click **Next**. In the Network Type window, verify that **All network connections** is selected, and click **Next**. In the IP Filter List window, click **Add**. In the Name text box, type **Domain Traffic**, and click **Add**.
10. In the welcome window of the IP Filter Wizard, click **Next**. In the IP Filter Description and Mirrored property window, type **All IP** in the Description text box, verify that the Mirrored box is checked, and click **Next**. In the IP Traffic Source window, click the **Source address** list box to see what types of sources can be specified, return the selection to **Any IP Address**, and click **Next**. In the IP Traffic Destination window, verify that **Any IP Address** is selected, and click **Next**. Because you are selecting all traffic rather than traffic for a specific protocol in the IP Protocol Type window, click **Next** and click **Finish**. Your IP filter list should look like the one in Figure 9-13. Click **OK**. In the IP Filter List window, click the **Domain Traffic** option button to enable it, and click **Next**.

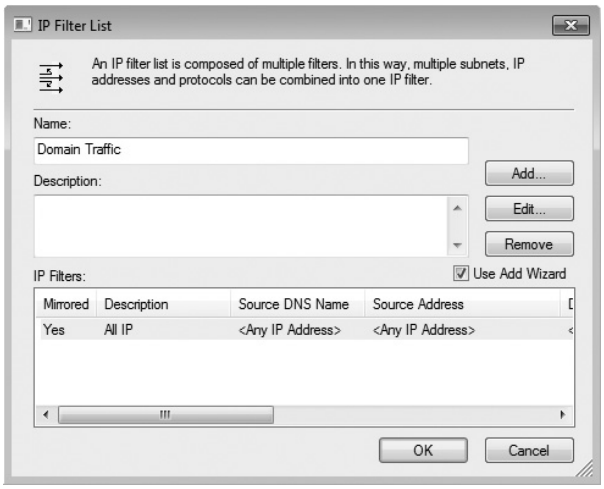

Used with permission of Microsoft Corporation

Figure 9-13 IP Filter List window

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

11. In the Filter Action window, you specify what happens when the target traffic is detected. Click **Add**. In the welcome window of the Filter Action Wizard, click **Next**. In the Filter Action Name window, type **Negotiate** in the Name text box, and click **Next**. In the Filter Action General Options window, verify that **Negotiate security** is selected, and click **Next**. In the next window, verify that **Do not allow unsecured communication** is selected, and click **Next**. In the IP Traffic Security window, verify that **Integrity** and **encryption** are selected, and click **Next**. Click **Finish**. In the Filter Action window, click the **Negotiate** option button, and click **Next**. In the Authentication Method window, you have the option of choosing Kerberos (Active Directory), digital certificates, or a passphrase (the least secure option). Leave Active Directory selected, and click **Next**. Click **Finish**.
12. In the New Rule Properties window, you can now see tabs for the different parameters you have configured. Click **Edit** in the various tabs to see the specific settings. Review the tabs and their settings to help you understand the process of setting up an IPsec rule. Click **OK** to close the New Rule window. You have now created an IPsec rule called Domain Traffic; if necessary, check its box to make it active. The Default response rule (<Dynamic>) should not be active. See Figure 9-14. Click **OK** to return to the Local Security Policy window, which should show your new Domain Connectivity Policy. Note that it is not assigned. Right-click the policy and click **Assign**. A green circle should appear on the policy's icon. See Figure 9-15.

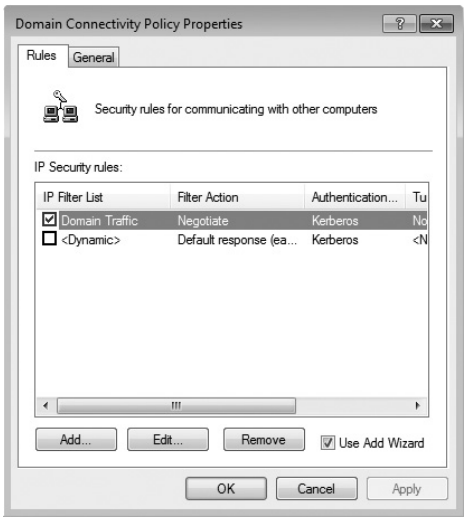

Used with permission of Microsoft Corporation

Figure 9-14 The Domain Traffic rule is enabled

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

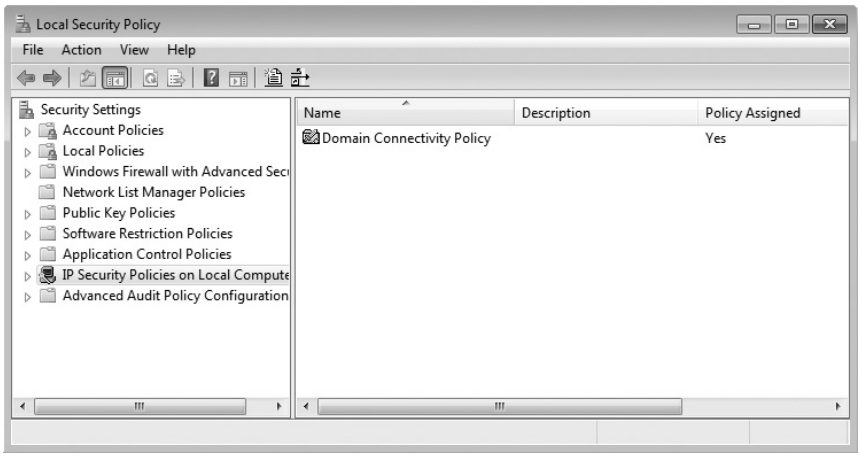

Used with permission of Microsoft Corporation

Figure 9-15 The Domain Connectivity Policy is enabled

13. Close all windows. Repeat Steps 2 and 3 to create a different file on the administrator's desktop on the Windows Server 2008 system. This time, the IPsec policy on Windows 7 complies with the policies on Windows Server 2008 and the connection is successful.
14. On Windows Server 2008, delete the Domain Authenticated Connections Policy from the Connection Security Rules section of the Windows Firewall with Advanced Security window.
15. On Windows 7, delete the Domain Connectivity Policy from the Local Security Policy's IP Security Policies on Local Computer section.
16. Leave the systems running for the next project.

Hands-On Project 9-3: Creating a Telnet Firewall Rule

Time Required: 30 minutes

Objective: Create a firewall rule to control Telnet traffic.

Description: Telnet is a protocol used for terminal emulation. Because it lacks support for encryption, Telnet is not used for sensitive communications across unsecure networks, but it makes a perfect protocol for demonstrating packet filtering. In this project, you enable Telnet on Windows Server 2008 and Windows 7, and then you learn to create firewall rules to control Telnet communications.

1. Log on to the Windows 7 system. Click **Start**. In the Search programs and files box, type **cmd**, and press **Enter**.
2. At the command prompt, type **telnet**, type the IP address of your Windows Server 2008 system, and press **Enter**. You receive an error message. Because Telnet is not a secure protocol, it is not enabled on Windows systems by default.
3. Click **Start**, click **Control Panel**, and click **Programs and Features**. In the Programs and Features window, click **Turn Windows features on or off**. In the next window, click the

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Telnet Client box, and click **OK**. Repeat Step 2. This time a different error message appears because Windows Server 2008 is not running the Telnet server service.

4. On Windows Server 2008, open **Server Manager**, click **Features**, click the **Add Features** link in the right pane, check the **Telnet Server** box, and click **Next**. Click **Install**. When the installation is complete, click **Close**.
5. From Windows 7, attempt Step 2 again. The command fails again.
6. From Windows Server 2008, open Server Manager, expand **Configuration**, click **Services**, and then scroll down and double-click **Telnet**. Again, for security reasons, Microsoft does not enable the Telnet Server service, even though it is installed. From the **Startup type** list box, select **Manual**. Click **Apply**, click **Start**, and click **OK**.
7. From Windows 7, attempt Step 2 again. The command succeeds, but Microsoft again wants to be sure that you are protected and asks you to confirm that you want to start the connection. Type **y** and press **Enter**. You are now managing the Windows Server 2008 system from the Windows 7 command prompt. You should see the C:\Users\Administrator prompt. If not, type **cd C:\Users\Administrator**, and press **Enter**. Type **cd Desktop** and press **Enter**. Type **dir** and press **Enter** to see the contents of the administrator's desktop. You should see the file called Private.txt that you created earlier in these projects. Type **exit** and press **Enter** to terminate the connection.
8. Now you will make a rule to block Telnet packets from entering Windows Server 2008. From the Windows Server 2008 system, open the Windows Firewall with Advanced Security window. Click **Inbound Rules** in the left frame. Click **New Rule** in the right frame. In the Rule Type window, click the **Port** option button, and click **Next**. In the Protocol and Ports window, verify that **TCP** is selected. In the **Specific local ports** text box, type **23**, which is the port used by the Telnet protocol. Click **Next**. In the Action window, click the **Block the connection** option button, and click **Next**. Uncheck the **Private** and **Public** boxes, leave the Domain box selected, and click **Next**. In the Name window, type **Telnet** in the Name text box, and click **Finish**. Note that the new rule, Telnet, is listed first in the Inbound Rules frame. The rule is shown with an icon of a red circle and slash, indicating that the rule blocks access.
9. Return to Windows 7, and attempt to establish a Telnet connection with the server again. The packets have been blocked by the server, so the command fails.
10. Log out of both systems.

Hands-On Project 9-4: Designing a Rule Base

Time Required: 30 minutes

Objective: Create a basic rule base for packet filtering.

Description: In this activity, you configure a set of basic packet-filtering rules for a network. The internal network is represented by 192.168.120.0; the firewall is hosted at 192.168.120.1; the e-mail server is at 192.168.120.2; the Web server is at 192.168.120.3; and the DNS server is at 192.168.120.4. Design the rule base using Table 9-13 to accomplish the following goals:

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

1. Create a rule that allows internal hosts to access the external network.
2. Create a rule that prevents access to the firewall.
3. Create a rule that allows internal and external access to the e-mail server and Web server.
4. Create a rule that allows internal access to the DNS server.

Rule	Source IP	Source port	Destination IP	Destination port	Action
1					
2					
3					
4					
5					
6					
7					
8					

© Cengage Learning 2014

Table 9-13 Rule base for packet filtering

Case Projects

Case Project 9-1: Improving a Rule Base

You have accepted a position as network administrator for a law firm. Unfortunately, the previous administrator did not leave much documentation of the network infrastructure and its configuration. You have spent many hours documenting the transmission media, data storage, and device configuration.

You have examined the firewall in an attempt to understand the existing rules and their purpose, and to see whether rewriting the rule base could speed up the firewall's performance and improve security. Table 9-14 shows the notes you have made about the firewall rule base. The firewall works correctly, but some simple changes to the rule base could improve its performance dramatically.

The network has IP addresses from 210.100.101.0 to 210.100.101.255. The firewall is at 210.100.101.1, the Web server is at 210.100.101.2, the DNS server is at 210.100.101.3, the SMTP server is at 210.100.101.4, and the POP3 server is at 210.100.101.5.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Rule	Source IP	Destination IP	Protocol	Action	Track	Comments
1	Any	210.100.101.1	Any	Deny	Alert	Blocks access to firewall
2	210.100.101.0 to 210.100.101.255	Any	S-HTTP	Deny	None	Blocks network access to Web server using S-HTTP
3	210.100.101.0 to 210.100.101.255	Any	HTTP, S-HTTP	Allow	None	Allows network access to all Web sites
4	Any	210.100.101.2	HTTP	Allow	Log	Allows all computers to access the Web server using HTTP
5	210.100.101.0 to 210.100.101.255	210.100.101.3	UDP	Allow	Log	Enables network to make queries to DNS server
6	210.100.101.3	Any except 210.100.101.0 to 210.100.101.255	TCP	Allow	Log	Enables DNS server to make lookups on the Internet but not in the network
7	210.100.101.0 to 210.100.101.255	210.100.101.5	TCP	Allow	None	Allows network access to POP3 server
8	Any	210.100.101.4	TCP	Allow	None	Allows any computer to access the SMTP server
9	Any	Any	Any	Deny	Log	Cleanup rule

© Cengage Learning 2014

Table 9-14 Firewall rule base

You have noted some questions that you need to address as you consider modifying the firewall rule base:

- Which rules cover the same sort of communication?
- Which rules are too far down the list and should be moved up?
- Which rules give the firewall more work than necessary? (*Hint:* Look in the Track column.)

On a separate piece of paper, create a rule base table. Using as few rows as possible, write a new rule base that addresses the questions in the preceding steps.

Case Project 9-2: Recommending a Software Firewall

Now that you have completed the new rule base for the law firm, as described in Case Project 9-1, your next task is to recommend a software firewall to install on the firm's workstations. Your supervisor wants you to consider both commercial and free products. Research both kinds of software firewalls, choose one product from each category, and then prepare a report that presents your choices and explains the advantages and disadvantages of each product.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 10

Firewall Design and Management

After reading this chapter and completing the exercises, you will be able to:

- Design common firewall configurations
- Describe, install, and configure a proxy server
- Describe and plan the configuration of a bastion host
- Describe and plan the configuration of a honeypot
- Describe and configure Network Address Translation
- Perform basic configuration of a Cisco ASA firewall

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

You can incorporate firewalls into your security infrastructure in many different ways. You can use one or more firewalls to provide protection for publicly accessible Web or FTP servers, protect file servers from external attack while allowing internal hosts access, and more. In fact, you can design your security architecture to fit any requirements you have. In this chapter, you learn about common firewall configurations and learn how to select one that is best suited to your needs.

You also learn about the functions of devices such as proxy servers, bastion hosts, honeypots, and firewalls. When deploying a defense-in-depth environment, the devices that host security functions must be as secure as possible themselves.

Designing Firewall Configurations

So far in this book, you have learned about routing, router security, and firewalls. In this chapter, you see how these components are combined on the network perimeter to protect the internal network and protect services in the demilitarized zone (DMZ). Businesses need a variety of services on their networks and on the Internet. They also need a network area where resource access is more flexible, allowing internal hosts and external systems to gain access. The problem is providing adequate access without jeopardizing confidential or mission-critical areas. To accomplish this goal, perimeter networks host Web servers, e-commerce servers, FTP servers, remote access services, databases, DNS servers, and more. Securing these systems requires firewalls, but firewalls cannot do the job alone.

To design a productive yet secure perimeter network, you must understand how firewalls and other security measures are combined with bastion hosts, Network Address Translation (NAT), proxy servers, and authentication methods. In the following sections, you learn how to deploy firewalls in different ways on a network: as part of a screening router, a dual-homed host, a screened host, a screened subnet DMZ, multiple DMZs, multiple firewalls, and a reverse firewall setup. A secure perimeter network can also include Internet Protocol Security (IPsec), virtual private networks (VPNs), and intrusion detection and prevention systems (IDPSs). The network might also provide wireless access.

Screening Routers

The simplest firewall is a single router on the network perimeter configured to filter packets. This **screening router** determines whether to allow or deny packets based on their source and destination IP addresses or other information in their headers (see Figure 10-1).

However, this device alone does not stop many attacks, especially those that use spoofed or manipulated IP address information. This type of router should be combined with a firewall or proxy server for added protection.

Dual-Homed Hosts

A common security arrangement is to install a firewall or other security software (such as a proxy server) on a **dual-homed host**, a computer configured with more than one network interface. The capability to forward packets is disabled on the computer, so only the firewall

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

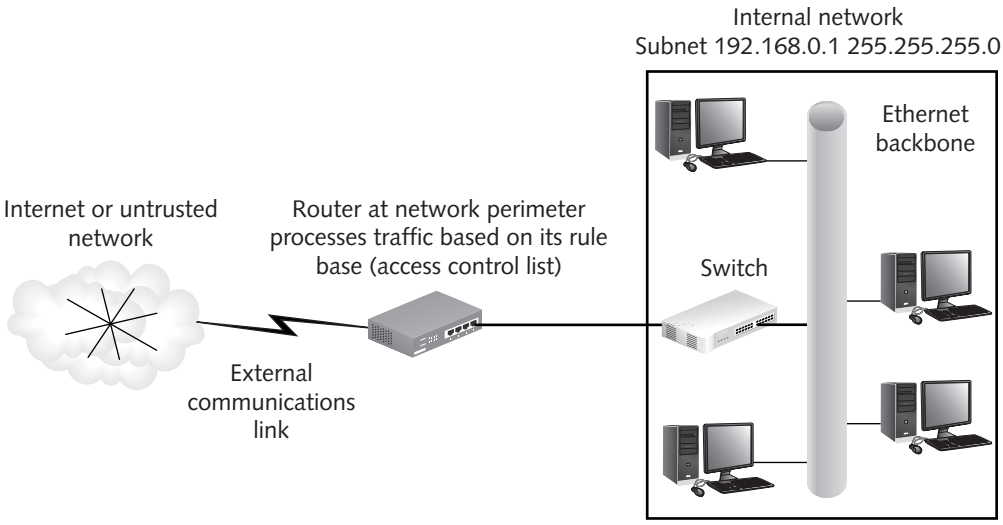

© Cengage Learning 2014

Figure 10-1 A screening router

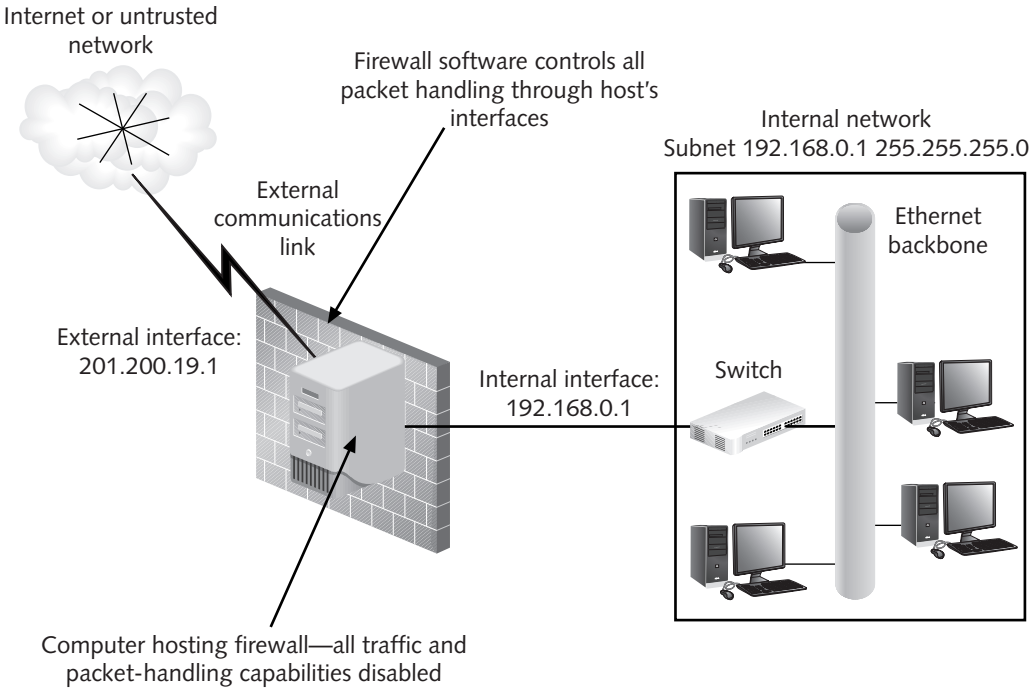

© Cengage Learning 2014

Figure 10-2 A dual-homed host

software can forward traffic from one interface to another (see Figure 10-2). Users who are connected to the Internet also use this type of setup. Rules are established so that the firewall moves traffic between the Internet and a home computer or network.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Originally, the term *dual-homed host* was used for a computer equipped with two separate network interface cards (NICs), with one NIC for each interface. Now this term is used to describe the setup in Figure 10-2, in which a firewall is placed between the network and the Internet. A dual-homed host provides only limited security because the firewall depends on the same computer used for day-to-day communication. Therefore, any problem with the host computer weakens the firewall. The big disadvantage is that the host serves as a single point of entry to the network, so attackers have to break through only one layer of protection to infiltrate the network. Therefore, a multilayered defense-in-depth (DiD) arrangement is even more important with a dual-homed host.

The computer that hosts the firewall software might have more than two network interfaces. It might be connected to the DMZ, the Internet, and the internal network, for example. In that case, it is called a *multihomed host*.

Screened Hosts

A **screened host** setup is similar to a dual-homed host, but the main difference is that a router is often added between the host and the Internet to carry out IP packet filtering.

A screened host is essentially a combination of a dual-homed host and screening router configuration; the two are blended for an added layer of functionality, security, and performance (see Figure 10-3). You might choose this setup for perimeter security on a corporate network,

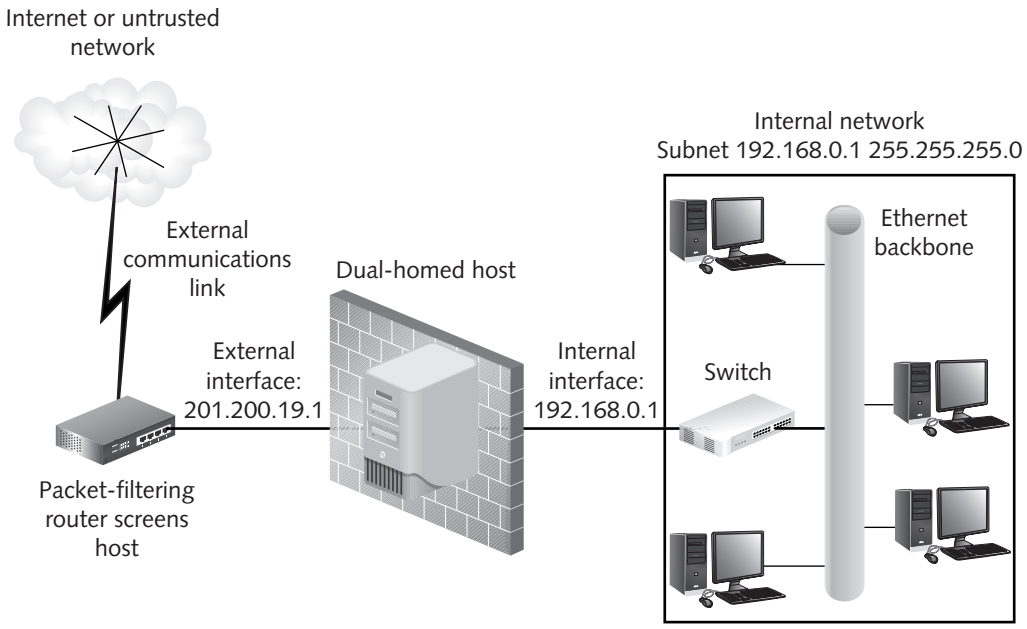

© Cengage Learning 2014

Figure 10-3 A screened host

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

for instance. A common enhancement is to have a screened host function as an application gateway or proxy server. The only network services allowed to pass through the proxy server are those for which proxy applications are already available.

Screened Subnet DMZs

A DMZ is a subnet of publicly accessible servers placed outside the internal network. Because this subnet contains important information, a packet filter or other security software should screen the DMZ. A common solution is to make the servers a subnet of the firewall (see Figure 10-4).

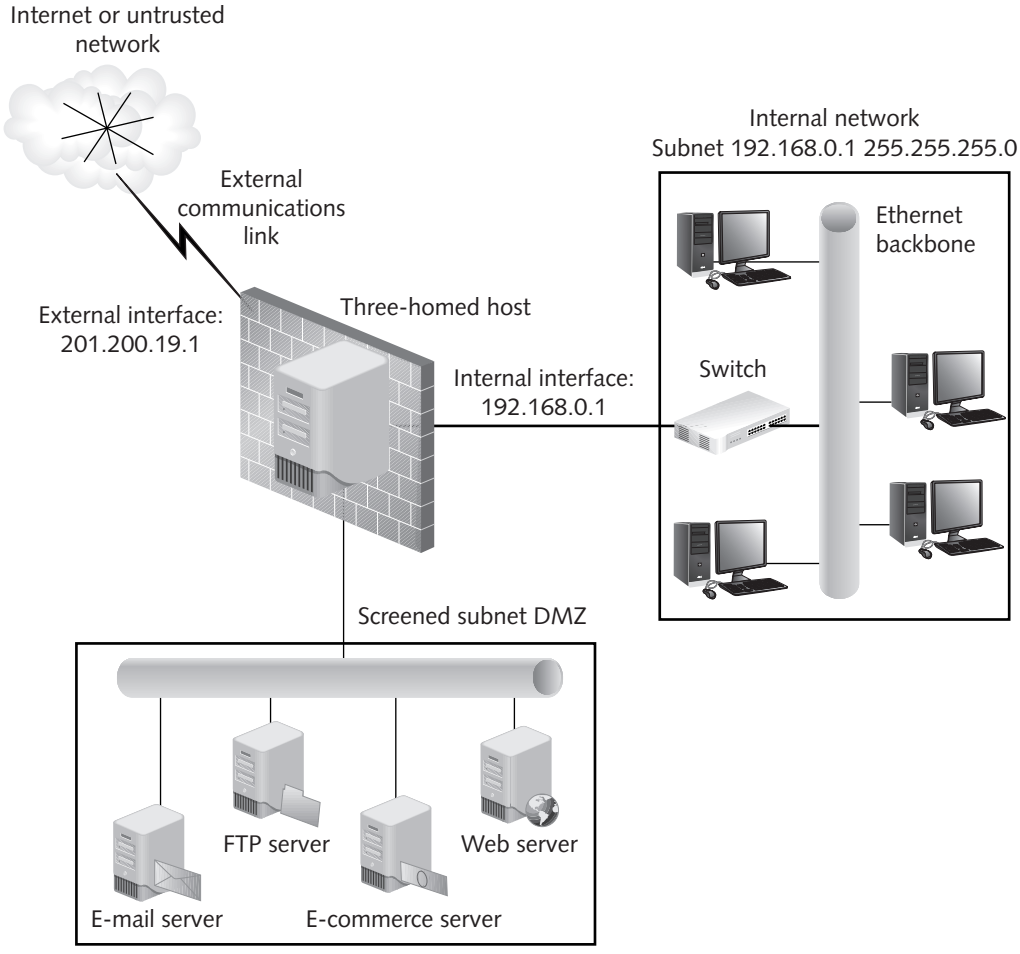

© Cengage Learning 2014

Figure 10-4 A screened subnet DMZ

Because the firewall that protects the DMZ is also connected to the Internet and can be connected to the internal network, it is often called a **three-pronged firewall**. You might choose this setup when you need to provide services to the public, such as an FTP server, a Web

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

server, or an e-mail server, but you want to make sure attackers cannot access your Web site or FTP resources.

The subnet that is attached to the firewall and contained in the DMZ is sometimes called a service network or perimeter network by those who dislike the military connotation of the term *DMZ*.

Multiple DMZ/Firewall Configurations

One DMZ does not provide enough security for many large corporations that are connected to the Internet or do business online. To handle demand from the Internet and keep the response time of public servers as fast as possible, you should set up multiple DMZs. Each DMZ is a **server farm**, a group of servers connected in a subnet that work together to receive requests with the help of **load-balancing software**. This software prioritizes and schedules requests and then distributes them to servers based on each server's current load and processing power.

Installing clusters of servers in DMZs outside the internal network helps protect the internal network from becoming overloaded. Placing the company's Web server in the outer subnet means it is better equipped to handle heavy traffic because there is less filtering of packets. If a Web server behind the firewall gets as many as 20,000 hits a minute, that amount of traffic could crash the server or seriously slow the firewall's performance as well as other traffic that needs to pass through. If the Web server is outside the protected network but in the DMZ, the firewall's performance is not compromised. To protect the local network, inbound connectivity from the Web server to the internal network should be blocked, and access attempts should be logged carefully.

Each server farm/DMZ should be protected by its own firewall or packet-filtering router. One possible configuration is shown in Figure 10-5. The service networks do not necessarily need to be protected from one another; instead, they feed into a single router, which sends traffic to the internal network. The service subnets are not connected directly to the Internet; a packet-filtering router screens each one.

A service network has also been established to allow sharing resources with a business partner. This network configuration should be handled with care, however, because some information could be sensitive; in this situation, placing a firewall between the packet-filtering router and business partner access subnet might be wise. A screened DMZ is also adequate, as long as the screening router is configured carefully, strict authentication methods are used, and access controls are placed on information at the file and folder levels. Encrypting sensitive data adds another layer of protection.

Finally, an additional firewall is placed inside the internal network to protect network management systems. The advantage of locating management software out of band—outside the internal network on a protected subnet of its own—is that management servers get an extra level of protection against intruders. If intruders gain control of the firewall or network management software, they might be able to access hosts on the network.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

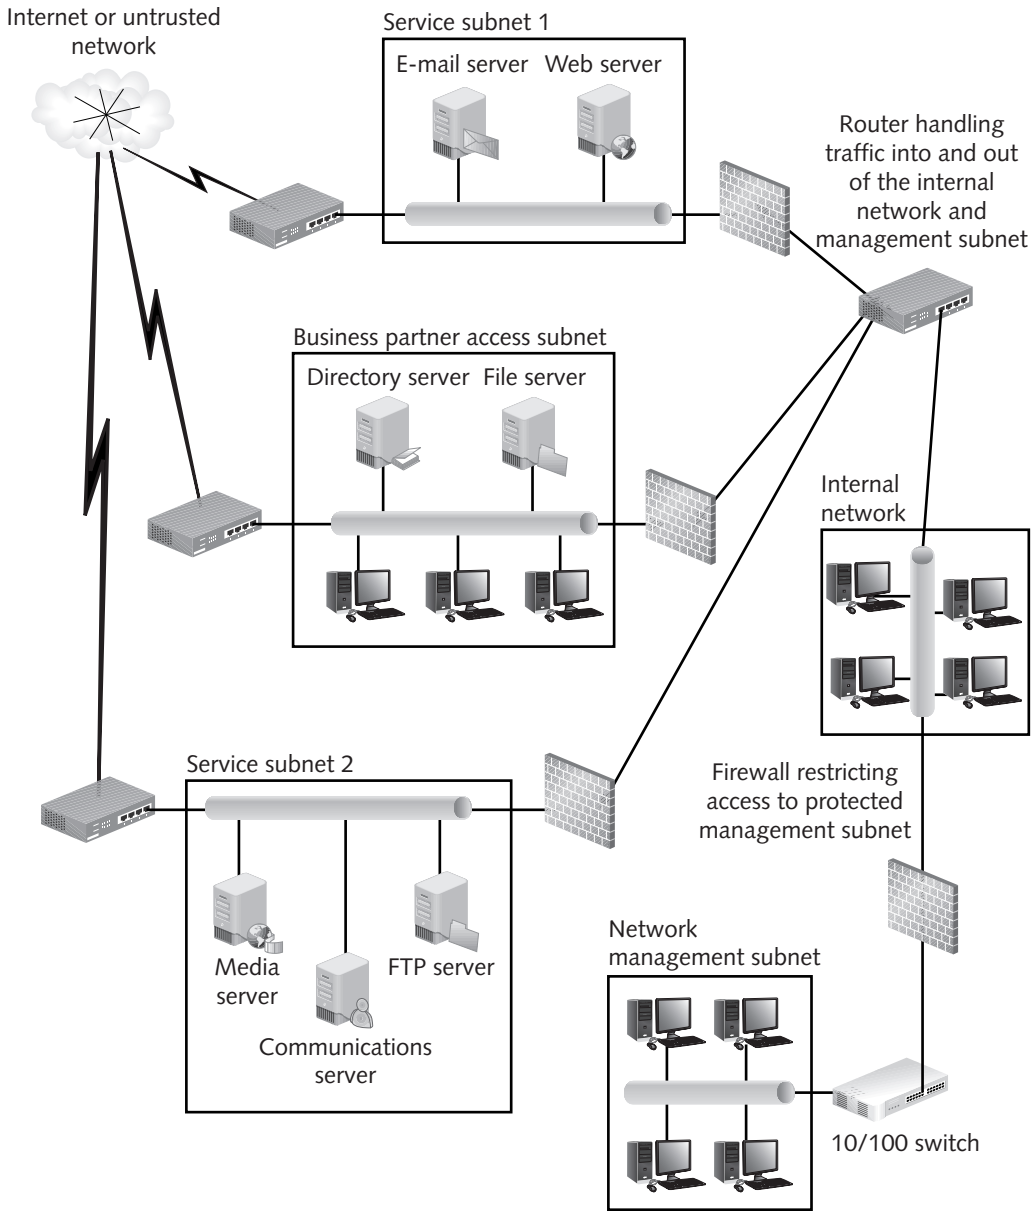

© Cengage Learning 2014

Figure 10-5 Multiple DMZs protected by multiple firewalls

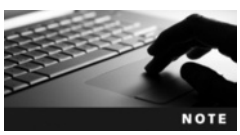

A protected subnet in an already protected internal network could be used as an extra layer of protection for confidential information, such as customer or personnel data. The extra layer might protect servers from disgruntled employees within the organization as well as attackers via the Internet.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Multiple Firewall Configurations

DiD means that multiple security devices are configured to work together to provide protection. DiD makes use of many different layers of network security. To achieve this level of protection, many organizations find that they need more than one firewall, either throughout the network or on a subnet. The following sections describe how to use two or more firewalls to protect an internal network, one DMZ, two DMZs, and branch offices that need to connect to the main office's internal network. In addition, multiple firewalls can help you achieve load distribution that keeps heavy traffic flowing through the gateway smoothly.

Protecting a DMZ with Multiple Firewalls When multiple firewalls are deployed around the network perimeter, they can work together to balance the network's traffic load. When multiple firewalls are used together, they must be configured identically and use the same firewall software. That way, traffic coming from the Internet can be balanced between the firewalls by using routers or switches on either side (see Figure 10-6).

The two firewalls in Figure 10-6 have an external interface on the Internet. You can also have one firewall with an interface on the Internet and a second firewall with an interface on the internal network; the DMZ being protected is positioned between them. Using two firewalls helps in the following ways:

- One firewall can control traffic between the DMZ and the Internet, and the other can control traffic between the protected network and the DMZ.
- The second firewall can serve as a **failover firewall**, which is a backup that can be configured to switch on if the first one fails, thus ensuring uninterrupted service.

A major advantage of setting up a DMZ with multiple firewalls is that you can control traffic in the three networks you are dealing with: the external network outside the DMZ, the external network within the DMZ, and the internal network behind the DMZ. You can identify certain protocols, such as outbound HTTP port 80, that should go to the external network within the DMZ, and you can allow other protocols to pass through to the internal network.

Using multiple interior routers to connect your DMZ subnet to parts of your internal subnets can cause problems. For example, using Routing Information Protocol (RIP) on an internal system might determine that the most direct route to another internal system is through the DMZ. As a result, confidential internal traffic flows across your DMZ, where it can be intercepted if an attacker manages to break into a host computer. Having multiple interior routers also makes overall network configuration more difficult.

Protecting Branch Offices with Multiple Firewalls A multinational corporation that needs to share information among branch offices in different locations can communicate securely by using a single security policy implemented by multiple firewalls. The main office has its own centralized firewall, which directs traffic for branch offices and their firewalls. The main office develops the security policy and deploys it through the centralized firewall on a dedicated computer called a **security workstation** (see Figure 10-7). Each branch office has its own firewall, but the main office develops and controls the security policy. The policy is then copied to other firewalls in the corporation.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

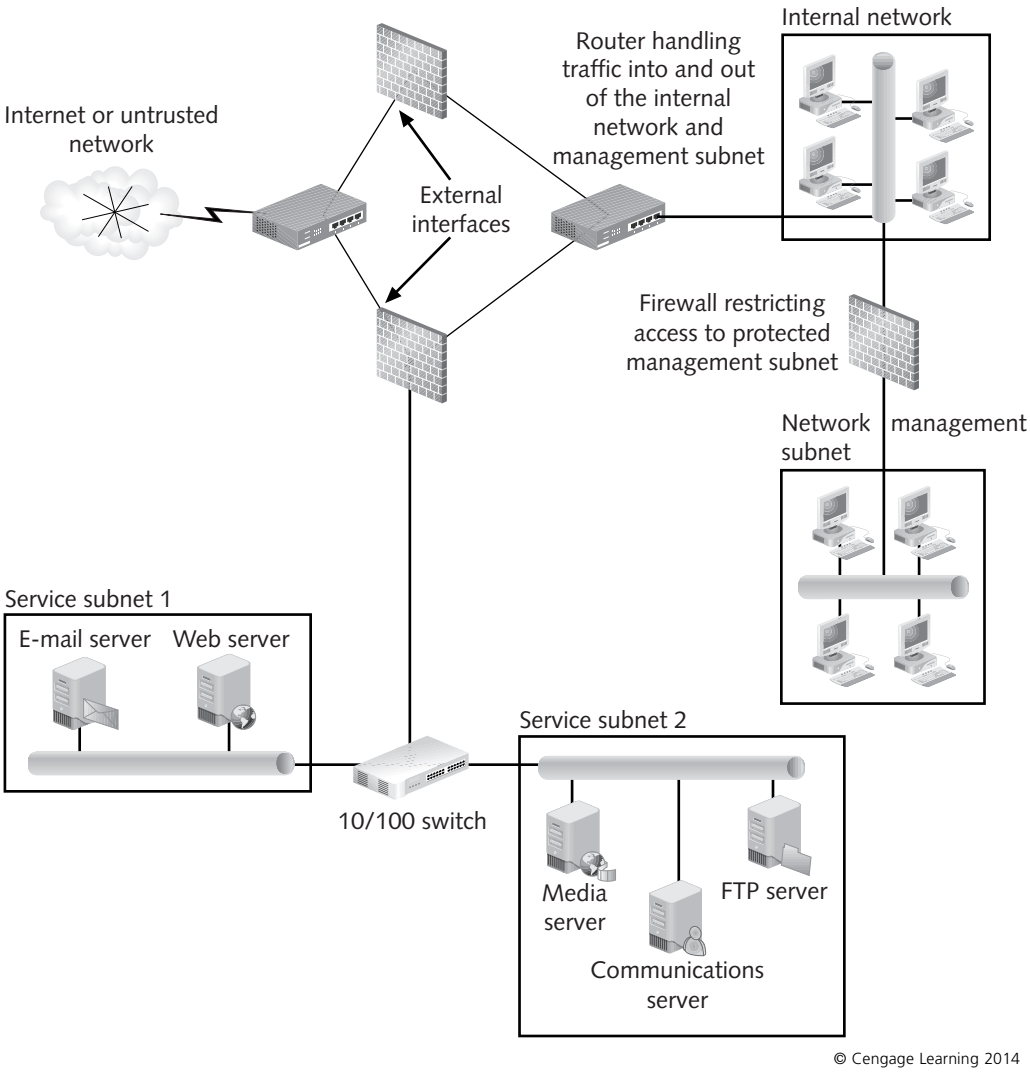

© Cengage Learning 2014

Figure 10-6 Two firewalls used for load balancing

One notable aspect of the setup in Figure 10-7 is that the two firewalls have a path for communicating with one another and a router to direct traffic to each firewall. However, traffic from the Internet or the corporate network does not travel between the two firewalls; only traffic from the security workstation travels on the connection between DMZs, carrying configuration information to firewalls and receiving log file data from them.

Reverse Firewalls

Some forward-thinking companies install a **reverse firewall**, a device that monitors outgoing connections instead of trying to block incoming traffic. Sometimes the biggest threats to an internal network come from its own users. For example, a university might make certain applications available only to researchers and scientists who need them. Clever computer

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

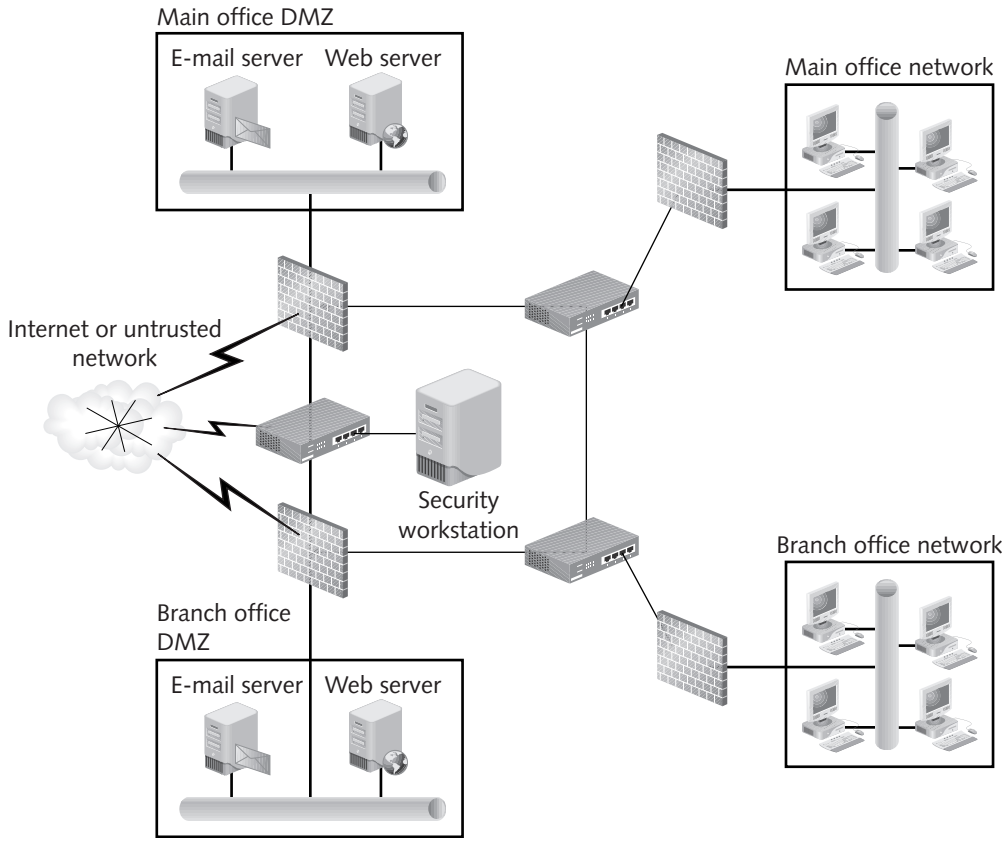

© Cengage Learning 2014

Figure 10-7 Multiple firewalls protecting branch offices

science students might try to break into the server that stores these applications so they can download the applications and use them. A reverse firewall would help by monitoring outgoing connection attempts that originate from internal users and then filtering out unauthorized attempts. Reverse firewalls have other purposes as well. For example, a company that is concerned about how its employees use the Web and other Internet services can set up a reverse firewall to log connections to Web sites. It can then block sites that are accessed repeatedly but are considered unsuitable for employees to visit during work hours.

In a DoS or distributed denial of service (DDoS) attack, information floods out of a network from compromised computers (commonly called zombies), thus overloading the network. A reverse firewall inspects outgoing packets and tracks their origin in the network. If the firewall detects many unexpected packets leaving the network, it sends a notification to the network administrator. This feature, however, can be programmed into any hardware or software firewall; it does not require a specialized reverse firewall.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Choosing a Firewall Configuration

Which firewall configuration is the best one for your needs? Many modern products have an advantage in the amount of customization available and their capability to perform multiple functions, such as firewall and IDPS services. Remember that firewall setups are not an either/or decision. You can combine a screened host with a multiple firewall setup, for instance. You can also have a reverse firewall in addition to a conventional firewall. Table 10-1 summarizes the advantages and disadvantages of each setup.

Configuration	Advantages	Disadvantages
Screening router	Simple, inexpensive; good for home applications if a stateful packet filter is used	Provides only minimal protection; viruses, Trojan programs, and some malformed packets might get through
Dual-homed host	Simple, economical; can provide effective protection if configured correctly	Provides a single point of entry (and fault); the firewall depends entirely on the host computer
Screened host	Provides two layers of protection for home and small-business networks	Provides a single point of entry (and fault); the firewall depends on the host computer and the router protecting it
Screened subnet DMZ	Protects public servers by isolating them from the internal network	Servers in the DMZ are highly vulnerable and need to be hardened
Multiple DMZs/firewalls	Provide layers of protection for a business network	Expensive
Single DMZ/two firewalls	Balance traffic load in high-traffic situations	Expensive
Branch offices/multiple firewalls	Provide protection for all offices in a corporate network as well as central administration	Firewalls must be purchased, installed, and configured at each office location
Reverse firewall	Monitors attacks from inside the network; enables organizations to monitor user activity	Can slow down user access to external networks or other parts of the internal network

© Cengage Learning 2014

Table 10-1 Advantages and disadvantages of firewall configurations

Examining Proxy Servers

Your organization's security policy might call for installing a **proxy server**—software that forwards packets to and from the network being protected and caches Web pages to speed up network performance. For many organizations, a proxy server is the only type of firewall they have installed, except for the company's router. The following sections describe the goals of setting up a proxy server and explain how it works, how to choose one, and how proxy servers can filter content for your network.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Proxy servers provide effective protection because they work at the Application layer of the OSI model. They can interpret what application is being used to make a request and forward the request on behalf of that application. In contrast, firewalls primarily interpret IP and TCP header information at the lower OSI levels.

Goals of Proxy Servers

Originally, speeding up network communications was the primary goal of proxy servers. The process, as shown in Figure 10-8, works like this: As Web pages are requested from a company's Web server, the proxy server receives the Web pages and forwards them to the computer making the request. At the same time, the proxy server caches the Web pages' text and image files—in other words, it stores those files on disk for later retrieval if needed. Computers that request the same Web page more than once have their requests received by the proxy server rather than the Web server. The proxy server checks its cache for the Web

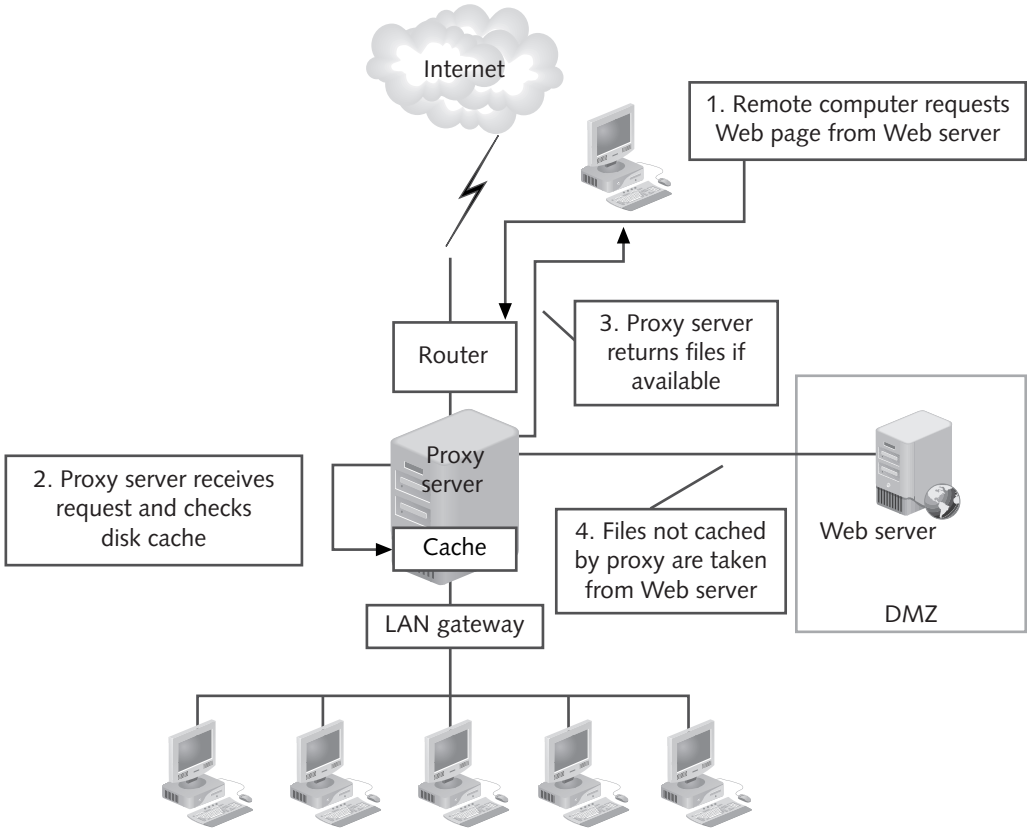

© Cengage Learning 2014

Figure 10-8 Proxy servers cache Web pages and other files

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

page being requested and compares the page's contents against those of the page currently published on the Web server. If no changes are found, the files are retrieved from the cache rather than from the original Web server. Storing documents in disk cache reduces the load on the Web server and speeds up network traffic.

The primary goal of modern proxy servers is to provide security at the Application layer and shield hosts on the internal network. A secondary goal is controlling which Web sites users are allowed to access. Proxy servers can use IP addresses or domain names to block access to specific Web sites or to entire top-level domains. For example, an administrator could allow access to the .gov top-level domain so that employees could view government forms online.

Many proxy servers come with default configurations that enable users to access the Internet using multiple services. However, these default configurations can open security holes. They might be set up to enable Telnet access, for example, or to enable Web access around the clock, which most users do not need. For better security, consider disabling services that most users do not need.

How Proxy Servers Work

Network and Transport layer packet filters look only at the header part of a TCP/IP packet. Their goal is to block unauthorized packets and allow only authorized packets to reach their destination. If packets are authorized, the packet filter enables host and client computers to communicate with one another directly.

In contrast, the goal of proxy servers is to prevent a direct connection between an external computer and an internal computer. One way proxy servers accomplish this objective is by working at the Application layer. When a request is received from an internal computer, the proxy server opens the packet and examines the data. For example, if the request is for a Web page and it uses the standard GET method, the proxy server reconstructs the packet and forwards it to the requested Web server, acting as a proxy Web browser. By acting at the Application layer, the proxy server can interpret which application was originally used to make a request and which application is needed to forward that request.

What does it mean to "reconstruct" a packet? When a proxy server receives a request, it opens the packet, examines the contents, and replaces the original header with a new header that contains the proxy server's own IP address instead of the original client's. An example is shown in Figure 10-9.

As you can see in Figure 10-9, the proxy server is in the DMZ. The following procedure occurs:

1. User 1 requests the Web page at IP address 188.92.101.5.
2. The proxy server receives a request from User 1's Web browser in the internal LAN, examines the request, and strips off the packet header. The proxy server replaces the header with its own public-source IP address before sending the packet on its way.
3. The Web server that receives the request interprets it as coming from the proxy server's IP address instead of from User 1's computer. In fact, the Web server (or an attacker who intercepts the request) has no way of knowing that User 1's computer exists.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

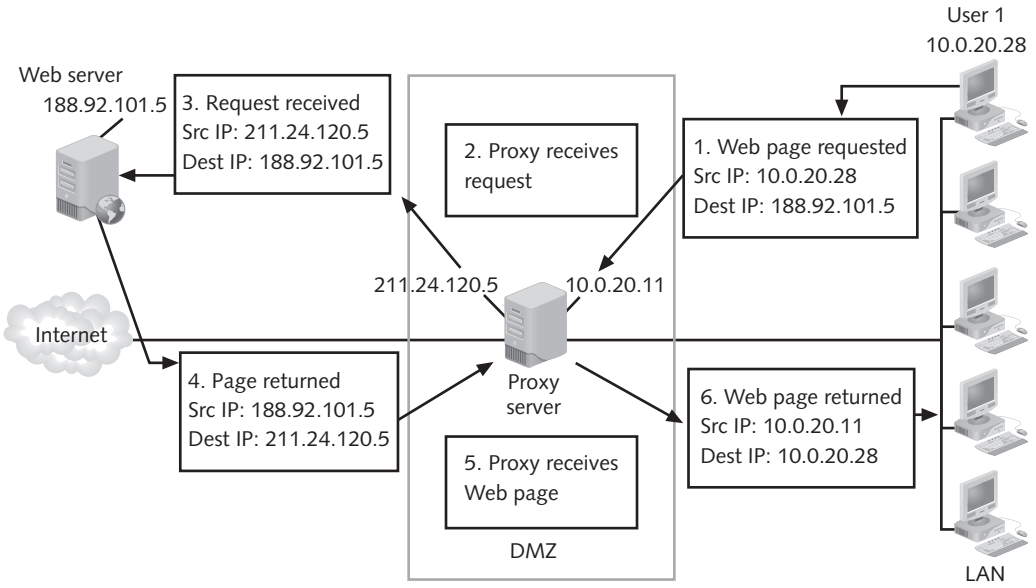

© Cengage Learning 2014

Figure 10-9 Proxy servers replace source IP addresses with their own addresses

4. The Web server then sends its response with the Web page to the proxy server.
5. The response goes through the proxy server for processing, and the IP header sent by the Web server is replaced.
6. The proxy server sends the requested Web page to User 1's computer, where the browser displays it.

The proxy server is configured to receive traffic before it goes to the Internet; client programs, such as Web browsers and e-mail applications, are configured to connect to the proxy server rather than the Internet. A typical browser configuration from Microsoft Internet Explorer is shown in Figure 10-10.

In Figure 10-10, the proxy server's IP address of 172.168.1.1 has been entered as the proxy address to use. Port 8080, which is normally used for proxy services, is also specified. This configuration results in the proxy server forwarding requests to external hosts over port 8080 while using a source IP address of 172.168.1.1. Proxy servers require all users on a network to configure client programs accordingly. Depending on the number of users in your network, configuration can be time consuming. In most networks, client settings such as proxy server addresses are configured automatically through network management.

Table 10-2 summarizes the advantages and disadvantages of using a proxy server.

Choosing a Proxy Server

Different proxy servers perform different functions to strengthen your existing firewall configuration. The type of proxy server you install depends on your network's needs. For example, you can install a simple freeware proxy server if you are satisfied with the level of

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

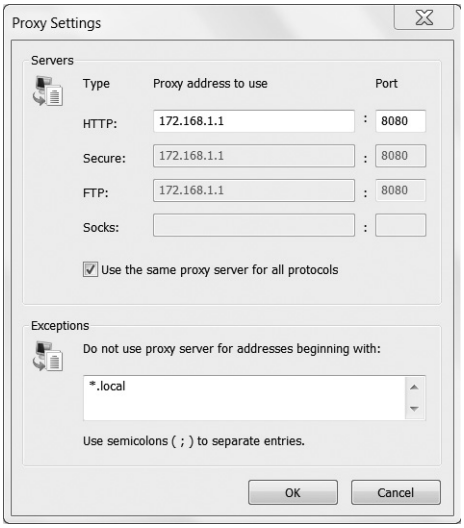

Used with permission of Microsoft Corporation

Figure 10-10 Configuring client programs to connect to the proxy server rather than the Internet

Advantages	Disadvantages
Examines contents of packets and filters on contents	Security can be weak depending on configuration
Shields internal host IP addresses	Can slow down network access
Caches Web pages for faster access	Might require configuration of client programs to use the proxy server
Provides a single point of logging	Provides a single point of failure

© Cengage Learning 2014

Table 10-2 Proxy server advantages and disadvantages

protection your existing firewall provides and you simply want to add functions the firewall cannot perform, such as filtering out pop-up ads, executable code, or other types of content.

On the other hand, if you want to strengthen existing firewall protection, a commercial proxy server can perform many functions that improve your overall network security, including a single point of logging and hiding internal IP addresses. If you need to save installation time and want to have only one program to manage, an enterprise-level firewall with a proxy server, packet filtering, and other functions is a good choice. The basic types of proxy servers you can choose—freeware proxy servers, commercial proxy servers, or a firewall that includes proxy server functions—are described in the following sections.

Freeware Proxy Servers Freeware proxy servers tend to offer a specific function rather than the full range of proxy server functions, so they are often described by names such as "content filter." An Internet search for freeware proxy servers returns dozens of links, but be aware that most of these products do not have the features needed for business

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

applications. One notable freeware proxy server is Squid for Linux (*www.squid-cache.org*). It is an open source product, so plenty of documentation and support are available.

Commercial Proxy Servers The many benefits of commercial proxy servers become evident when you install a program such as Microsoft Forefront Threat Management Gateway (TMG). This commercial program combines the capabilities to cache Web pages and translate source and destination IP addresses with content filtering and traditional firewall functions, such as packet filtering and NAT. TMG includes advanced features such as HTTPS inspection, which allows the proxy server to decrypt and inspect traffic that has been protected by asymmetric and symmetric encryption. To compete with other security programs, most proxy servers are not advertised as having these features. Instead, they are described as part of a more comprehensive firewall package. A commercial proxy server is a good choice for a business network if you plan to upgrade the software as new versions become available. Any commercial program should offer technical support to help with installation and configuration problems.

Proxy Servers That Can Include Firewall Functions Some proxy servers, such as TMG, can be set up to act as firewalls in addition to performing other duties. Having an all-in-one program simplifies installation, product updating, and day-to-day management. On the other hand, all your network security needs are being handled by a single program. If something goes wrong with your proxy server, your firewall also goes down. If feasible, using several software and hardware programs in a coordinated network defense layer is preferable. For example, you could use TMG for your proxy server needs and Cisco ASA as your firewall.

Filtering Content

A useful feature of proxy servers is their capability to open TCP/IP packets, inspect the data portion, and take action based on the contents. This capability enables proxy servers to filter out content that would otherwise appear in a user's Web browser. In a business environment, proxy servers can be configured to block Web sites that contain content employees should not be allowed to view. They can also drop any executable programs, such as Java applets or ActiveX controls, embedded in Web pages that can potentially damage files or replicate files when they run on a user's computer.

Choosing a Bastion Host

Security software does not operate on its own. It is installed on a computer that runs on an operating system (OS); this computer needs to be as secure as possible because of the important software it contains and its position on the network perimeter. Often, this computer also functions as a server that provides Web pages, e-mail, and other services to users inside and outside the network. This practice is discouraged because of vulnerabilities caused by the complexity of configuring a system for multiple services; however, small companies might not have the resources to host these services separately. In this situation, a company should consider hiring a security specialist.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

To protect security software as well as the network, the computer should be turned into a **bastion host**, a computer on the network perimeter that has been specially protected with OS patches, authentication, and encryption. A computer hosting a firewall that is riddled with OS vulnerabilities gives attackers a way to compromise the entire network. They might even be able to disable firewall software if they can gain administrative privileges or crack passwords on the server. The following sections describe how to choose and configure a bastion host that can safely host security software.

General Requirements

In general, a bastion host should be a computer running an OS that is already secure or has been in release long enough that patches for security vulnerabilities are available. When the OS is made as secure as possible by eliminating unnecessary software and services, closing potential openings, and protecting information with encryption and authentication, the computer is said to be hardened.

The following steps outline the general process of creating a bastion host. These steps are covered in more detail in the following sections.

1. Select a machine with adequate memory and processor speed.
2. Choose and install the OS and any patches or updates.
3. Determine where the bastion host fits in the network configuration. Make sure it is in a safe and controlled physical environment.
4. Install the services you want to provide, or modify existing services.
5. Remove services and accounts that are not needed.
6. Back up the system and its data, including log files.
7. Conduct a security audit.
8. Connect the system to the network.

Selecting the Bastion Host Machine

You do not need to select the latest hardware and software combination to configure a bastion host. Instead, choose a combination of machine type and software that you are familiar with and can work with easily. You do not want to be repairing or rebuilding a machine under pressure and learning to operate it at the same time.

In an ideal situation, you can designate one host for each service you want to provide: one FTP server/bastion host, one Web server/bastion host, one SMTP server/bastion host, and so on. However, budget constraints might force you to combine services on one bastion host. In this case, a comprehensive risk analysis of the services and hardware you need to protect the most can be helpful. The following sections discuss essential components of a secure bastion host.

Choosing an Operating System

The most important requirements for a bastion host are your level of comfort with the OS and its inherent security and reliability. Management probably is not concerned about which

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

OS you install; your priorities are ensuring that the machine protects the internal network and being certain you can get it running and maintain it easily.

Another top priority is the security level of the OS. Make sure you select a version that is stable and secure, and check the OS's Web site or well-known security sites for patches and updates for new vulnerabilities. Some popular security sites are noted in the following list:

- Windows Server 2008 (*www.microsoft.com/en-us/server-cloud/windows-server/default.aspx*)
- Red Hat Linux (*www.redhat.com*)
- FreeBSD Project (*www.freebsd.org*)
- SANS Institute's list of the Top Cyber Security Risks, which includes a subsection on Application vs. Operating System Patching; the SANS Web site is at *www.sans.org*
- U.S. Department of Energy's Joint Cybersecurity Coordination Center (JC3) site (*www.doecirc.energy.gov*), which lists newly discovered security advisories

Memory and Processor Speed

Bastion hosts do not need to have the most recent or expensive memory-processor combinations. Memory is always important when operating a server, but because the bastion host might be providing only one service, you are not likely to need many gigabytes of RAM. However, you might need to match processing power to the server load, which could mean simply upgrading the processor or perhaps adding processors.

Location on the Network

Bastion hosts are usually located outside the internal network and are often combined with packet-filtering devices, such as routers and firewall appliances, on either side (see Figure 10-11). This combination further protects the bastion host from attack because it can prevent tampered-with packets from ever reaching the bastion host. In addition, a packet-filtering server (proxy server) can filter out suspicious packets that come from inside the network as a result of Trojan programs or viruses.

More often, multiple bastion hosts are set up in the DMZ, with each machine providing a single service to the public (see Figure 10-12). This configuration is more secure because you can allow general access to the Web server yet limit mail server access to employees, for example.

Hardening the Bastion Host

A bastion host can be any server that hosts Web, e-mail, FTP, and other network services. However, the more services you install on the bastion host, the greater the chance of introducing a security vulnerability on the system. Therefore, one way to harden a bastion host is by removing all unnecessary software, services, and user accounts. The simpler your bastion host is, the easier it is to secure. In addition, services could have software bugs or configuration errors, which could lead to security problems. To give intruders fewer opportunities to exploit the system, only the minimum number of services and open ports should be available.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

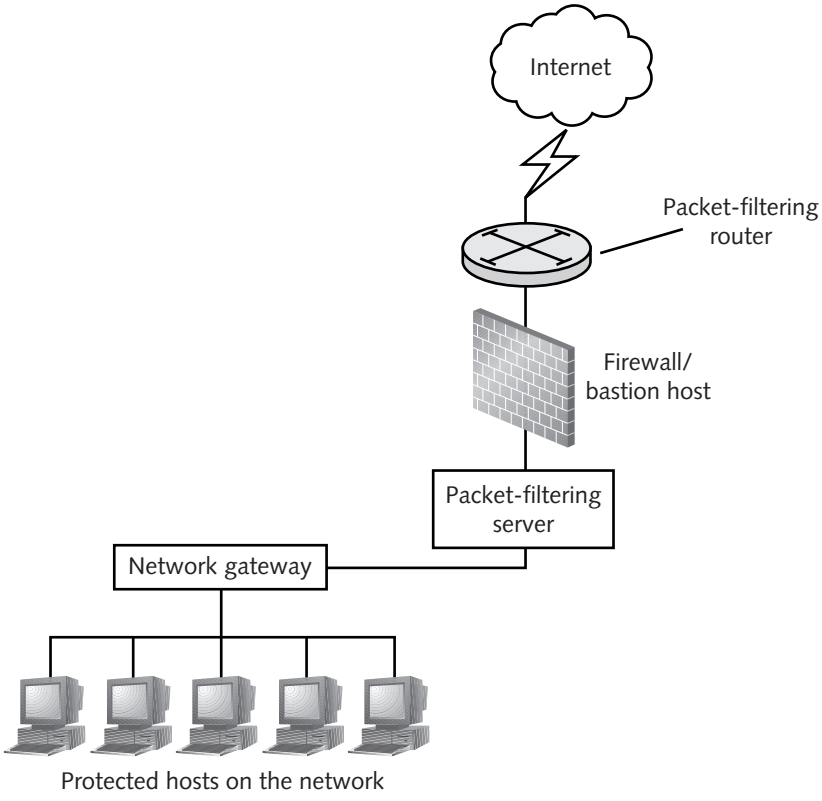

© Cengage Learning 2014

Figure 10-11 Bastion hosts are often combined with packet-filtering routers

Selecting Services to Provide

A bare-bones configuration reduces the risk of attacks and has the extra benefit of boosting efficiency. If you close unnecessary ports and disable user accounts and services you do not plan to use, attackers have fewer ways to access the system. Some critical services you should disable are the ones that perform routing or IP forwarding—unless, of course, the bastion host is intended to function as a router. Disabling IP forwarding makes it more difficult for attackers to communicate with internal network computers.

When you are stopping or removing services, you should not disable any dependency services—services the system needs to function correctly. For example, Telnet is usually disabled because other services do not normally require it. Remote Procedure Calls (RPC), however, should not be disabled because almost everything depends on it to function correctly. In addition, you should stop services one at a time to see what effect each has on the system. Get in the habit of documenting every change you make and how the system reacts to the change; this record helps if you need to troubleshoot later. This documentation process, called change management, could be added as a requirement of your security policy. Backing up the system before making major changes to services is also a wise precaution.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

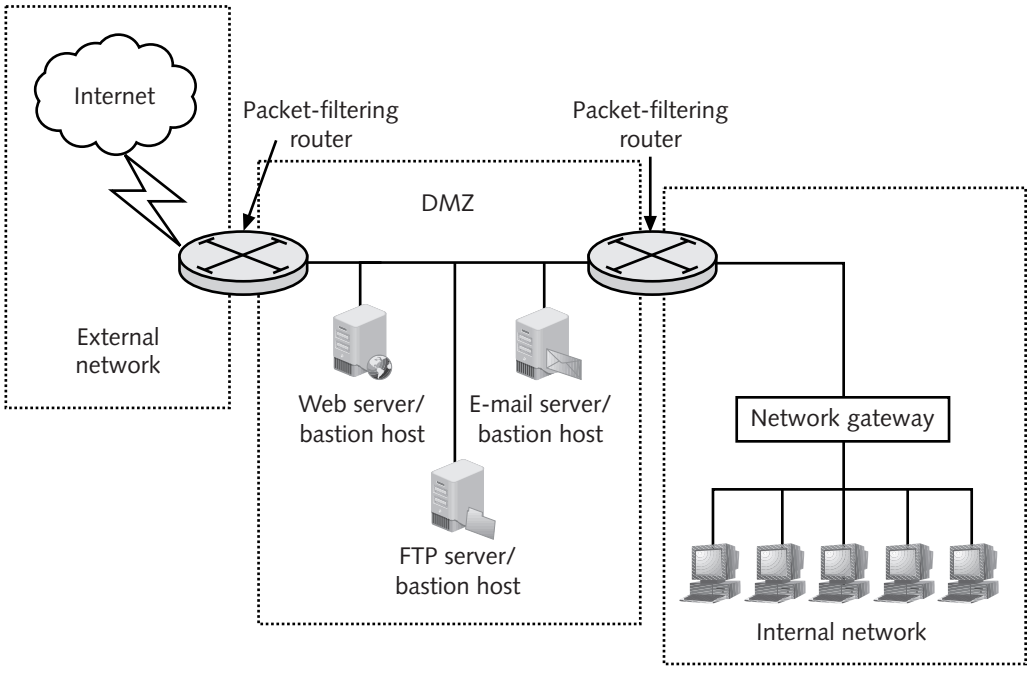

© Cengage Learning 2014

Figure 10-12 Bastion hosts in the DMZ

Using Honeypots

A honeypot is a computer placed on the network perimeter to attract attackers so that they stay away from critical servers on the network (see Figure 10-13); it may or may not be a bastion host. This computer is equipped with software and perhaps data files that appear to be important to the company. A honeypot might also be configured with security holes so that it seems vulnerable to known attacks. A honeypot can be located between the bastion host and internal network; if an attacker manages to get past the external packet filter to the DMZ and then scan for open ports, the attacker could get "stuck" in the honeypot—in other words, diverted from your real files by being misdirected to files of no value.

Network security experts are divided over the use of honeypots. Some think they have value, some think they are outdated and not worth using, and some think they are unnecessary and even potentially dangerous if they contain information about your company and its bastion hosts. Laws concerning the use of honeypots are confusing at best, so consult your legal department before setting one up. Honeypots are worth study not only because they are likely to be covered on certification exams, but because they are still discussed as options in perimeter security.

Another goal of a honeypot is logging. Because intruders who cannot tell the difference between the honeypot and a legitimate target can be expected to attack the honeypot, you can configure it to log every access attempt as a way to identify attackers. Because your

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

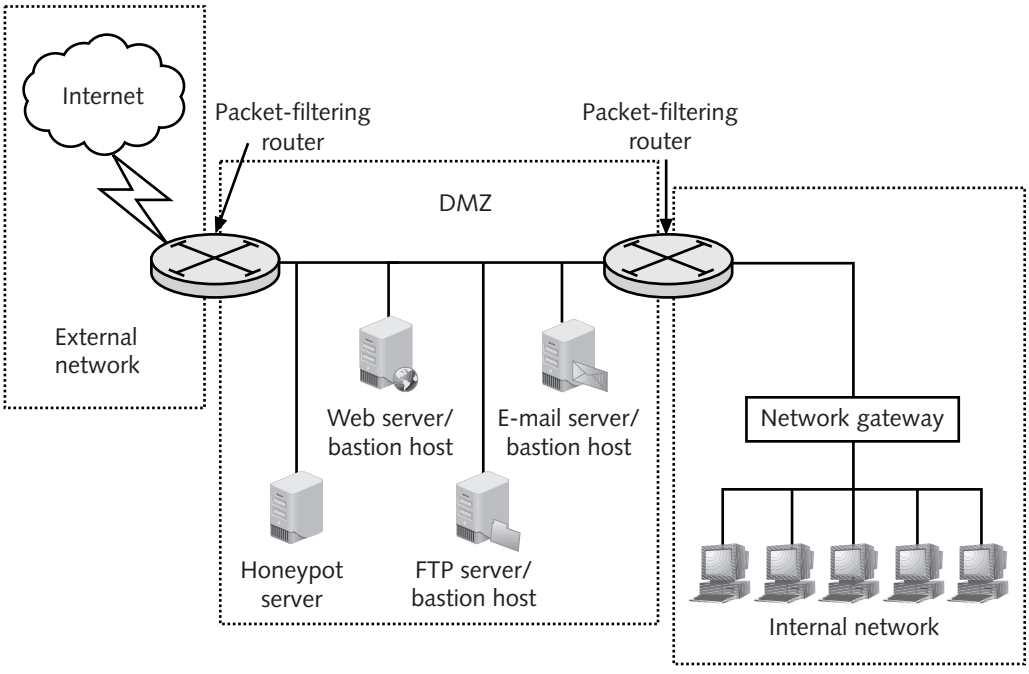

© Cengage Learning 2014

Figure 10-13 A honeypot in the DMZ

employees are not accessing resources in the honeypot, any access attempt is probably from an attacker. A honeypot can also provide indications of how your real network computers might be attacked. By reviewing what OS flaws, open ports, or other vulnerabilities are exploited on the honeypot, you can take steps to address these vulnerabilities on other systems.

Disabling User Accounts

Default accounts are sometimes created during installation of OSs and other software, and some of these accounts have blank or default passwords. You should delete or disable all user accounts from the bastion host. They are not needed because users should not be able to connect to the bastion host from their computers. User accounts on the bastion host increase the chances of a security breach.

You should also rename the Administrator account as another way to thwart intruders. Many attackers can access computers through administrative accounts that use the default name "Administrator" and are never assigned a password. Renaming these accounts and using long, complex passwords can help prevent attacks.

You should assume that the bastion host will be compromised in some way, and then take proactive steps to secure it and prepare for associated security incidents. The bastion host is the system most likely to be attacked because of its availability to external users.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Handling Backups and Auditing

Setting up a bastion host requires a systematic and thorough approach; backups, detailed recordkeeping, and auditing are essential steps in hardening a computer.

Bastion hosts can generate a lot of log files and other data, such as alert messages. You need to copy this information to other computers in your network regularly. When you perform detailed recordkeeping, the information goes through network defense layers you have already set up. Because of its high security configuration, your bastion host will probably be in a vulnerable location on the DMZ and outside the internal network. Log files and system data, which need to be backed up regularly, should go through the firewall that protects the internal network to screen them for viruses and other vulnerabilities, such as manipulated packets. Attackers might insert falsified information into packets as a way to access computers on the internal network.

Auditing should be configured for all failed and successful attempts to log on to the bastion host and for any attempts to access or change files. Administrators rarely have enough time to review log files, but you should make time every day to review logs for any devices in the DMZ as well as routers and packet-filtering servers. These logs can give you advance warning of intrusion attempts.

To avoid using network resources, be sure that systems used as bastion hosts have DVD-RW drives, removable disk drives, or tape drives so that you can make backups.

Network Address Translation

Originally, **Network Address Translation** (**NAT**) was designed to help conserve the store of public IP addresses. Private, internal networks did not need to use public IP addresses as long as the internal clients could access outside network hosts through a public IP address on a NAT device in their network. See Figure 10-14.

This technology allowed administrators to assign private IP address ranges in the internal network. Because these addresses are not routable on the Internet, NAT provided security and conserved the public IP address space. By shielding IP addresses of internal hosts, a NAT device makes it more difficult for attackers to find computers to exploit. Many attacks begin with an intruder locating a computer with a static public IP address; the intruder can then scan the computer for open ports to exploit. If the intruder cannot find the computer's IP address, an attack might not be possible.

The NAT process begins with an internal client sending a packet destined for an external host to the NAT device. The NAT device repackages the internal client's packet and sends it to the destination host on the Internet. When the Internet host responds, the NAT device receives the packets, and then refers to a table in its memory to determine which internal client should get the packet.

This address mapping process typically occurs in two ways; the terminology associated with these two address translation types has long been a source of confusion. For the purposes of

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

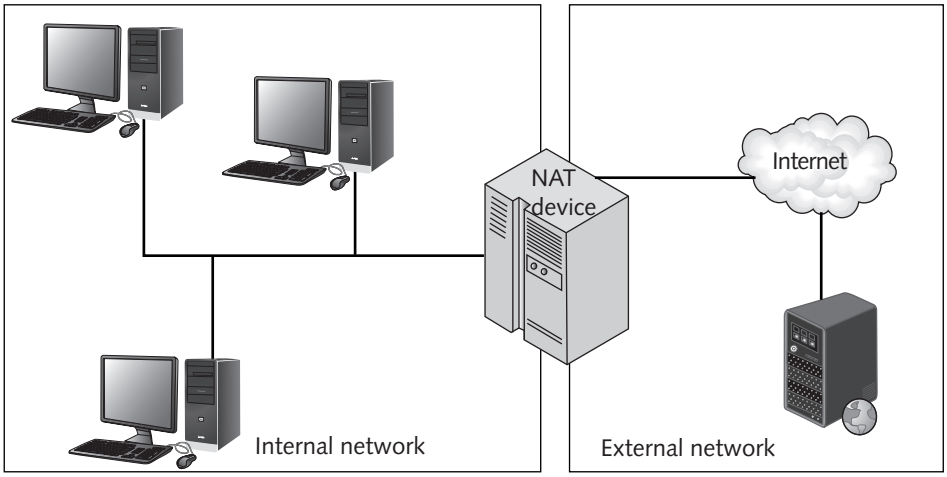

© Cengage Learning 2014

Figure 10-14 Internal clients communicate with external hosts through a NAT device

clarity, *Network Address Translation* is the term used in this book to apply to both primary address translation types: one-to-one NAT and many-to-one NAT.

One-to-One NAT

One-to-one NAT is the process of mapping one internal IP address to one external IP address. This process is summarized in the following list and shown in Figure 10-15.

1. The internal client at 172.16.25.12 sends packets destined for the Internet Web server at 206.188.95.117 to its default gateway on the NAT device at 172.16.25.1.
2. The NAT device repackages the packet so that its public interface, 97.205.16.211, appears to be the source of the packet, and then sends the packet on to the Web server.
3. The Web server responds to the NAT device.
4. The NAT device repackages the Web server's response and sends it to the internal client.

Static mapping is a version of one-to-one NAT that is often used when private addressing is used on a DMZ and all traffic for a busy DMZ server, such as a Web server, needs to have a dedicated public IP interface. In this case, the Web server's address is statically assigned, and all traffic sent to the specified public NAT IP address is sent to the static IP address of the Web server.

Many-to-One NAT

If one-to-one NAT were the only type of address translation, there would have to be one public IP address on a NAT device for each internal client that wanted simultaneous access to the Internet. Obviously, this would be expensive and would not conserve public IP addresses. **Many-to-one NAT**, sometimes called Port Address Translation, uses TCP or UDP

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

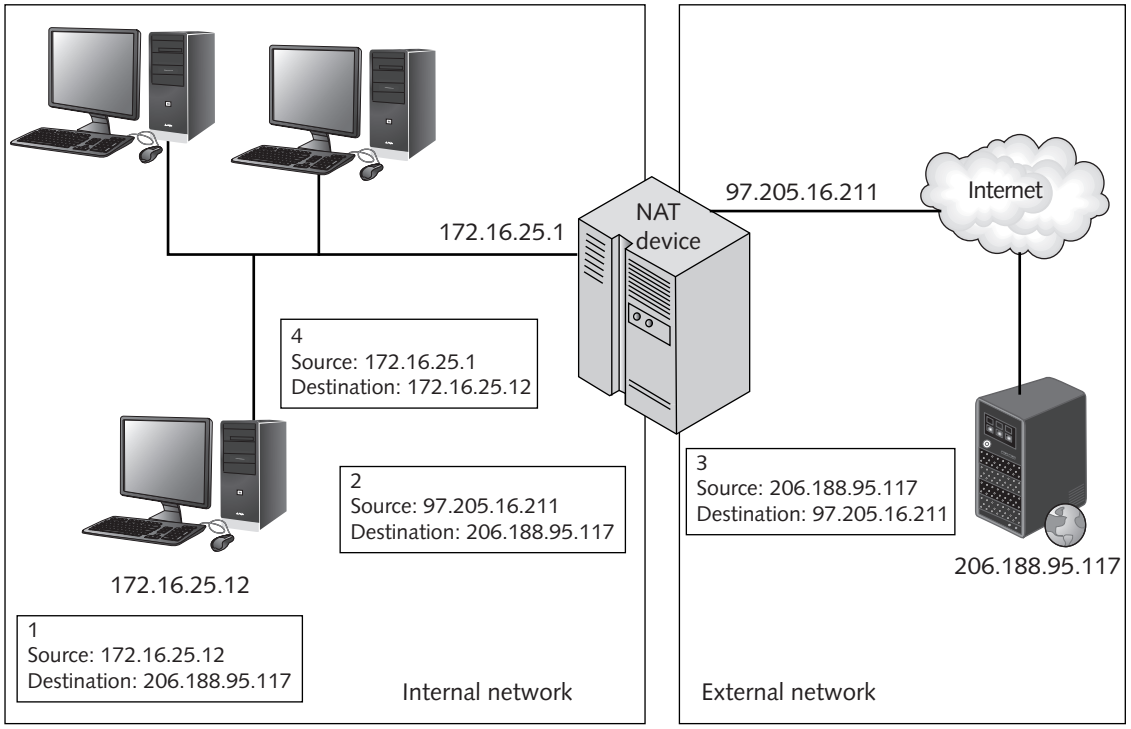

© Cengage Learning 2014

Figure 10-15 One-to-one NAT

port addresses to distinguish between internal clients, which allows many internal clients to use the same single public NAT interface simultaneously. As in one-to-one NAT, computers on the Internet see only the NAT device's IP address; packets from internal hosts seem to be coming from the NAT device. This mode has disadvantages, however. First, you can hide only so many clients behind a single IP address. For most networks, this limitation is not a big problem; for a large network, however, performance can degrade as the number of connections increases. In addition, many-to-one NAT does not work with some types of VPNs because the two endpoints must have unique addressing. Finally, many-to-one NAT uses only a single public IP address, so you cannot provide other services, such as a Web server, unless you have another IP address for them.

In Figure 10-16, the address translation also includes Port Address Translation as follows:

1. The internal client at 172.16.25.12:2250 sends packets destined for the Internet Web server at 206.188.95.117:80 to its default gateway on the NAT device at 172.16.25.1. (Note that these addresses use the socket format *IP:port*.)
2. The NAT device repackages the packet so that its public interface, 97.205.16.211:1788, appears to be the source of the packet, and then sends the packet on to the Web server. The NAT device receives the packet at its port 1788 and keeps a record of which internal host should get the packet.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

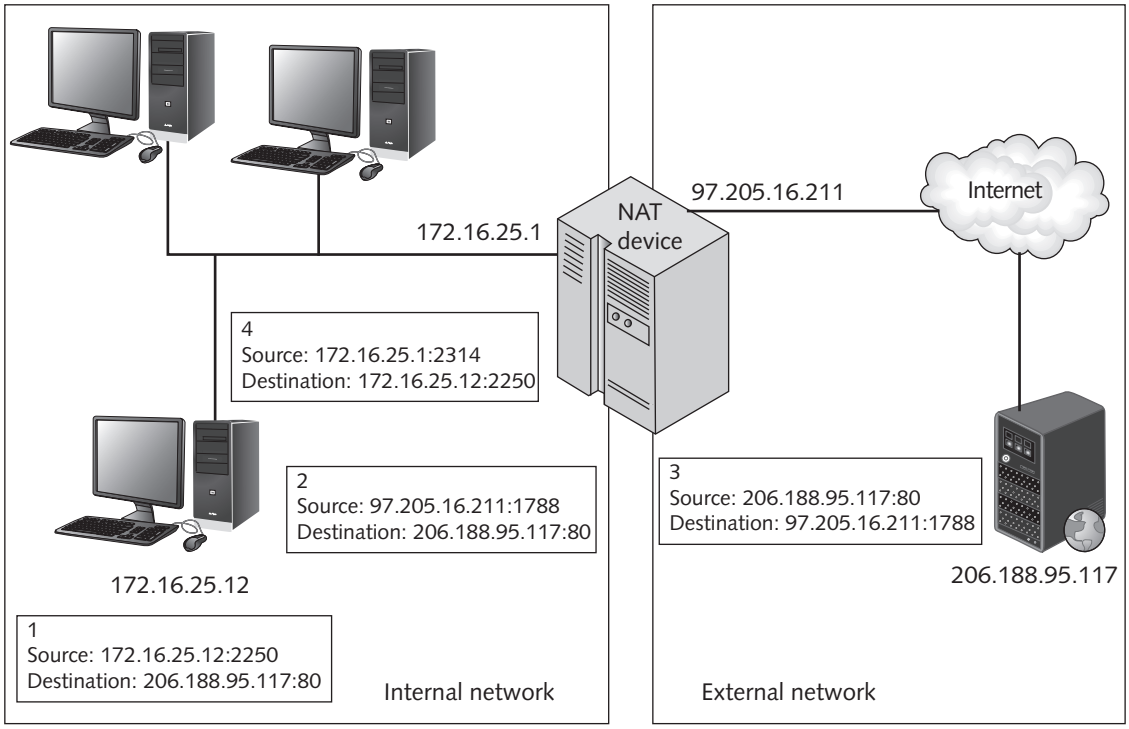

© Cengage Learning 2014

Figure 10-16 Many-to-one NAT

3. The Web server responds to the NAT device.
4. The NAT device consults its table, determines which host should receive the packet, repackages the Web server's response, and sends it to the internal client.

A single NAT device essentially sets up a firewall for an internal network. In fact, any security scheme that shields computers from external attack is considered a firewall, regardless of whether the hardware or software device is labeled as one.

Because NAT devices function like a go-between for internal and external computers, they can be confused with proxy servers. Both devices shield internal hosts by forwarding requests from the Internet to computers and vice versa. However, proxy servers rebuild packets before sending them on, and NAT simply repackages the packets.

Firewall Configuration Example

Different hardware firewall products exist to meet a wide range of customer needs from small home networks to massive networks with tens of thousands of connections. Most enterprise devices can be configured either from a command line or from a Web-based graphical user interface, although typically a command-line interface is used for system configuration.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

In this section, you learn the basics of configuring a Cisco ASA (Adaptive Security Appliance) 5505 firewall, as shown in Figures 10-17 and 10-18. This product is appropriate for small businesses and branch offices. It supports SSL and IPsec VPNs and contains PoE (Power over Ethernet) ports to support VoIP (Voice over IP) phones. The ASA 5505 includes a built-in switch.

Courtesy of Cisco Systems, Inc. Unauthorized use not permitted

Figure 10-17 Front view of ASA 5505

Courtesy of Cisco Systems, Inc. Unauthorized use not permitted

Figure 10-18 Rear view of ASA 5505

Connecting to the ASA 5505 is the same as connecting to a Cisco router, as discussed in Chapter 4. A rollover cable is connected to the management PC's COM 1 port and to the firewall's Console port. A terminal emulator such as PuTTY is used to make the command-line connection.

The command prompt is "ciscoasa" by default, and the enable password is blank. Thus, by typing "enable" and then pressing Enter when prompted for the password, you are placed in privileged mode. The show switch vlan command demonstrates that all eight switch ports are placed in VLAN 1 by default:

```
ciscoasa# show switch vlan
VLAN Name                           Status       Ports
---- ------------------------- --------- -------------------------
 1                                 down     Et0/0, Et0/1, Et0/2, Et0/3
                                            Et0/4, Et0/5, Et0/6, Et0/7
```

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

By switching to global configuration mode, you can name the firewall:

```
ciscoasa# configure terminal
ciscoasa(config)# hostname SanFrancisco
SanFrancisco(config)#
```

At this point, it would be wise to assign a strong password to protect the enable command:

```
SanFrancisco(config)# enable password T%imPwa0)gi
```

If you wanted to make Ethernet 0 the connection to the Internet and Ethernet port 1 the DMZ while leaving the other Ethernet ports for the internal LAN, you would proceed as follows:

```
SanFrancisco(config)# interface ethernet 0/0
SanFrancisco(config-if)# no switchport access vlan 1
SanFrancisco(config-if)# switchport access vlan 2
SanFrancisco(config-if)# exit
SanFrancisco(config)# interface ethernet 0/1
SanFrancisco(config-if)# no switchport access vlan 1
SanFrancisco(config-if)# switchport access vlan 3
SanFrancisco(config-if)# exit
SanFrancisco(config)# exit
SanFrancisco# show switch vlan
VLAN Name                           Status    Ports
---- ------------------------------ --------- ---------------------------
1    -                              down      Et0/2, Et0/3, Et0/4, Et0/5
                                               Et0/6, Et0/7
2    -                              down      Et0/0
3    -                              down      Et0/1
```


In these commands, you have removed Ethernet ports 0 and 1 from VLAN 1 (the LAN) and added them to VLANs 2 and 3 (the Internet and the DMZ, respectively). The show switch vlan command verifies the configuration.

Now you can name the VLANs and assign them security levels and IP addresses:

```
SanFrancisco# configure terminal
SanFrancisco(config)# interface vlan 1
SanFrancisco(config-if)# nameif LAN
INFO: Security level for "LAN" set to 0 by default.
SanFrancisco(config-if)# security-level 100
SanFrancisco(config-if)# ip address 192.168.1.205 255.255.255.0
SanFrancisco(config-if)# exit
SanFrancisco(config)# interface vlan 2
SanFrancisco(config-if)# nameif INTERNET
INFO: Security level for "INTERNET" set to 0 by default.
SanFrancisco(config-if)# ip address 209.57.12.126 255.255.255.0
SanFrancisco(config-if)# exit
SanFrancisco(config)# int vlan 3
SanFrancisco(config-if)# no forward interface vlan 1
```

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

```
SanFrancisco(config-if)# nameif DMZ
INFO: Security level for "DMZ" set to 0 by default.
SanFrancisco(config-if)# security-level 50
SanFrancisco(config-if)# ip address 10.0.0.1 255.255.255.0
SanFrancisco(config-if)# exit
SanFrancisco(config)# exit
```

In these commands, you have named the three zones LAN, INTERNET, and DMZ. Security levels are set automatically and range from least secure at 0 to most secure at 100. Thus, the LAN (the internal network) needed to have higher security and was set at 100, whereas the Internet zone security was set at 0. Notice that the DMZ was configured not to forward traffic to the internal LAN. Use the show ip address command to show the current configuration:

```
SanFrancisco# show ip address
System IP Addresses:
Interface              Name       IP address      Subnet mask     Method
Vlan1                  LAN        192.168.1.205   255.255.255.0   manual
Vlan2                  INTERNET   209.57.12.126   255.255.255.0   manual
Vlan3                  DMZ        10.0.0.1        255.255.255.0   manual
```

Now would be a good time to save your configuration to nonvolatile RAM so that the system will "remember" its settings if it is accidentally turned off. Enter the following command:

```
SanFrancisco# copy running-config startup-config
Source filename [running-config]? [Press Enter]
Cryptochecksum: 667668f0 e75003da 64fd6d58 aaa178f4

1539 bytes copied in 1.560 secs (1539 bytes/sec)
SanFrancisco#
```

If you have a TFTP server, you should save your configuration there as well so that if the whole router goes down, you can restore the configuration to a new one quickly:

```
SanFrancisco# copy startup-config tftp
Address or name of remote host []? 192.168.1.201
Destination filename [startup-config]? SanFran
!
1539 bytes copied in 2.700 secs (769 bytes/sec)
SanFrancisco#
```

You can also verify the IP interfaces:

```
SanFrancisco# show interface ip brief
Interface         IP-Address     OK? Method Status                Protocol
Internal-Data0/0  unassigned     YES unset  up                    up
Internal-Data0/1  unassigned     YES unset  administratively down
Loopback0         127.0.0.1      YES unset   up                    up
Vlan1             192.168.1.205  YES manual  up                    up
Vlan2             209.57.12.126  YES manual  up                    up
Vlan3             10.0.0.1       YES manual  up                    up
Ethernet0/0       unassigned     YES unset   up                    up
```

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

```
Ethernet0/1    unassigned   YES unset  up                    up
Ethernet0/2    unassigned   YES unset  up                    up
Ethernet0/3    unassigned   YES unset  administratively down down
Ethernet0/4    unassigned   YES unset  administratively down down
Ethernet0/5    unassigned   YES unset  administratively down down
Ethernet0/6    unassigned   YES unset  administratively down down
Ethernet0/7    unassigned   YES unset  administratively down down
```

Note that the Ethernet interfaces 3 through 7, which are assigned to the LAN, have no hosts attached. The only LAN host is attached to Ethernet interface 2.

To allow the firewall to route between interfaces, you must enable routing. In this case, you will use RIP as in Chapter 4. The following commands implement and verify RIP:

```
SanFrancisco# configure terminal
SanFrancisco(config)# router rip
SanFrancisco(config-router)# network 10.0.0.0
SanFrancisco(config-router)# network 209.57.12.0
SanFrancisco(config-router)# network 192.168.1.0
SanFrancisco(config-router)# version 2
SanFrancisco(config-router)# exit
SanFrancisco(config)# exit
SanFrancisco# show route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
           area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

C    127.0.0.0 255.255.255.0 is directly connected, _internal_loopback
C    192.168.1.0 255.255.255.0 is directly connected, LAN
C    10.0.0.0 255.255.255.0 is directly connected, DMZ
C    209.57.12.0 255.255.255.0 is directly connected, INTERNET
```

Access control lists work much like they do in Cisco routers, as discussed in Chapter 4. Depending on the services your network needs to support, you need to be sure that you allow the required traffic. In the following commands, you prevent LAN clients from connecting to a range of IP addresses to which you have traced a number of attacks (67.35.166.0/24), but you allow all other connections. The access control list, named HackSource, could be created and applied as follows:

```
SanFrancisco# configure terminal
SanFrancisco(config)# access-list HackSource extended deny tcp
192.168.1.0 255.255.255.0 67.35.166.0 255.255.255.0
```

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

```
SanFrancisco(config)# access-list HackSource extended permit ip
any any
SanFrancisco(config)# interface eth 0/0
SanFrancisco(config-if)# access-group HackSource in interface LAN
```

Other types of filtering are possible. For example, if you wanted to remove Java applets in traffic from a Web server at 198.55.164.79 and port 80 to your LAN, you could use the following command:

```
SanFrancisco(config)# filter java 80 192.168.1.0 255.255.255.0
198.55.164.79 255.255.255.0
```

Another traffic control feature is TCP normalization. Although this feature allows you to be very granular in choosing which TCP traffic to allow or deny, by default it drops or clears packets that are considered abnormal. You can use this feature to verify checksums, allow or drop packets that exceed the maximum TCP segment size, and allow or drop invalid ACK flags, among other things.

Chapter Summary

- Firewall design includes planning the most effective location for firewall placement. This location depends on the amount of traffic that must be filtered, the level of security needed, and the types of assets being protected.
- You can use multiple firewalls when you need multiple DMZs or to provide load balancing. Multiple firewalls are also useful for fault tolerance and when an organization has branch offices.
- A proxy server was originally designed to improve Web access performance by caching Web pages for multiple clients. It is now used to perform firewall and Network Address Translation (NAT) tasks as well as Web caching.
- Bastion hosts are computers such as Web servers, e-mail servers, and proxy servers that are accessible to untrusted clients. Bastion hosts should be configured for maximum security.
- NAT is used to protect internal clients from direct access by untrusted, external hosts and to decrease the need for public IP addresses. By using the private IP address ranges as specified in RFC 1918, organizations can purchase only the public interface needed to support a large number of freely usable private IP addresses. The two main types of NAT are one-to-one, which has a direct mapping between an internal and external interface, and many-to-one, in which a large number of private addresses are mapped to a single external interface. The latter is performed by mapping TCP and UDP port addresses on the source and destination fields in the packet headers.
- Many of the same commands used to configure Cisco routers and switches are also applicable on Cisco firewalls. For example, on the Cisco ASA 5505 firewall, VLANs are created to distinguish internal, external, and DMZ networks.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Key Terms

bastion hosts Computers that are accessible to untrusted hosts.

dual-homed host A computer configured with more than one network interface.

failover firewall A backup firewall that is configured to switch on if the current firewall fails.

load-balancing software Software that prioritizes and schedules requests and then distributes them to servers in a server cluster based on each server's current load and processing power.

many-to-one NAT A process that uses the source and destination TCP and UDP port addresses to map traffic between internal and external hosts. Many-to-one NAT is also called Port Address Translation.

Network Address Translation (NAT) The repackaging of packets so that internal IP addresses are stripped from requests to an untrusted network like the Internet.

one-to-one NAT The process of mapping one internal IP address to one external IP address.

proxy server Software that forwards packets to and from the network being protected and caches Web pages to speed up network performance.

reverse firewall A device that filters outgoing connections.

screened host A dual-homed host in which one interface is connected to an internal network and the other interface is connected to a router to an untrusted network.

screening router A router placed between an untrusted network and an internal network.

security workstation A computer dedicated to providing firewall policies.

server farm A group of servers connected in a subnet that work together to receive requests.

three-pronged firewall A firewall with separate interfaces connected to an untrusted network, a semitrusted network, and a trusted network.

Review Questions

1. A DMZ is _____.
   a. a trusted network
   b. a semitrusted network
   c. an untrusted network
   d. not actually a network

2. A screening router would be an appropriate choice for meeting the security needs of a _____.
   a. small office network
   b. home network
   c. DMZ
   d. none of the above

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

3. Which of the following computers is likely to be found in a DMZ? (Choose all that apply.)
   a. e-mail server
   b. domain controller
   c. Web server
   d. customer information database
4. Which of the following issues should you consider in firewall design? (Choose all that apply.)
   a. fault tolerance
   b. log size
   c. authorization
   d. load balancing
5. A proxy server _____. (Choose all that apply.)
   a. is designed to improve Web access
   b. is the same as a reverse firewall
   c. uses fewer system resources than a software firewall
   d. can filter Application layer content
6. What is the main problem with using a screening router?
   a. The router can be configured incorrectly.
   b. The router might not provide an adequate screen.
   c. The router cannot be used with a firewall.
   d. The router alone cannot stop many types of attacks.
7. What enables servers in a server farm to work together to handle requests?
   a. a router
   b. a switch
   c. a networking hub
   d. load-balancing software
8. For which of the following reasons would you consider creating a protected subnet within an already protected internal network? (Choose all that apply.)
   a. to protect customer information
   b. to protect management servers
   c. to protect the company's reputation
   d. to protect Web servers

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

9. A corporation with several branch offices has decided to maintain multiple firewalls, one to protect each branch office's network. What is the most efficient way to maintain these firewalls?
   a. Use a centralized security workstation.
   b. Send information about the security policy to each network administrator.
   c. Set up remote desktop management software.
   d. Broadcast configuration instructions periodically by e-mail.
10. Which of the following functions can a bastion host perform? (Choose all that apply.)
   a. FTP server
   b. e-mail server
   c. security management server
   d. domain controller
11. Which of the following can hide internal IP addresses from the Internet? (Choose all that apply.)
   a. packet filters
   b. NAT
   c. proxy servers
   d. state tables
12. Hardening a bastion host involves which of the following measures? (Choose all that apply.)
   a. disabling unnecessary services
   b. removing unnecessary accounts
   c. installing current patches
   d. all of the above
13. To isolate all external Web requests to a specific Web server on the DMZ, it would be best to use many-to-one NAT. True or False?
14. A bastion host is usually located on the internal network. True or False?
15. In a Cisco ASA 5505 firewall, security level 100 is the least secure level. True or False?

Hands-On Projects

Hands-On Project 10-1: Installing Threat Management Gateway 2010

Time Required: 45 minutes

Objective: Install Threat Management Gateway 2010.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Description: Microsoft Forefront contains layers of security functionality, one of which is the proxy server/firewall called Threat Management Gateway 2010 (TMG). In this project, you install TMG as an edge device. Then, in later hands-on projects, you configure TMG to separate your internal network from a simulated Internet environment.

The TMG system must contain two NICs and support the 64-bit version of Windows Server 2008 R2.

The projects in this chapter are designed for four computers: the previously installed Windows Server 2008 domain controller, TMG running on another Windows Server 2008 system, the previously installed Windows 7 system, and the previously installed Linux system. IP addresses used in these hands-on projects are designed for one group of four computers. The instructor will assign each group a set of first octets that are different from those shown in these hands-on projects. For example, in the projects, the IP address ranges are 192.*x.y.z* and 10.*x.y.z*. The next group of four computers could be assigned IP address ranges of 193.*x.y.z* and 11.*x.y.z*, and so on.

1. Start your Windows Server 2008 domain controller. Set its IP address as 10.0.0.125 and its subnet mask as 255.0.0.0. Set its default gateway to 10.0.0.111 and its DNS server to 10.0.0.125.
2. Start your Windows 7 client system. Set its IP address to 10.0.0.110 and its subnet mask as 255.0.0.0. Set its DNS server to 10.0.0.125.
3. Install Windows Server 2008 R2 in the TMG system and configure it as follows:
   - Hostname: Team*x*TMG
   - Member of a domain: Team*x*.net
   - Install all patches and updates through Windows Update
   - Local Area Connection name: Inside
   - Local Area Connection 2 name: Outside
   - Inside IP address: 10.0.0.111
   - Inside subnet mask: 255.0.0.0
   - Outside IP address: 192.168.1.110
   - Outside subnet mask: 255. 255. 255.0
4. Start the Threat Management Gateway installation wizard. Click **Run Preparation Tool**. Click **Next**.
5. In the License Agreement window, click **I accept the terms of the License Agreements**, and click **Next**. In the Installation Type window, verify that **Forefront TMG services and Management** is selected, and click **Next**. In the Preparation Complete window, click **Finish**. The Forefront TMG Installation Wizard starts.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

6. In the welcome window, click **Next**. In the License Agreement window, click **I accept the terms in the license agreement**, and click **Next**. In the Customer Information window, type **Team*x*User** in the User Name text box, where *x* is the team number assigned by your instructor. In the Organization text box, type **Team*x***. Enter the product serial number if required, and click **Next**.
7. In the Installation Path window, note the installation path, and click **Next**.
8. In the Define Internal Network window, click **Add**. In the Addresses window, click **Add Adapter**. Click the **Inside** box, and click **OK**. (See Figure 10-19.) In the Addresses window, click **OK**. In the Define Internal Network window, click **Next**. Read the Services Warning window, and click **Next**. In the Ready to Install the Program window, click **Install**.

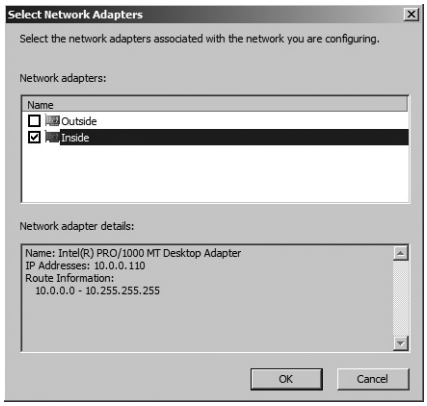

Used with permission of Microsoft Corporation

Figure 10-19 Selecting the inside adapter

9. In the Installation Wizard Completed window, click **Finish**.
10. Click **Start**, click **All Programs**, click **Microsoft Forefront TMG**, and click **Forefront TMG Management**.
11. In the Getting Started Wizard window, click **Configure network settings**. In the welcome window, click **Next**. In the Network Template Selection window, verify that the **Edge firewall** option is selected, and click **Next**.
12. In the Local Area Network (LAN) Settings window, select **Inside** from the Network adapter connected to the LAN list box. This selection should populate the IP address and Subnet mask text boxes. Click **Next**.
13. In the Internet Settings window, the settings should be populated automatically. Click **Next**, and click **Finish**.
14. In the Getting Started Wizard window, click **Configure system settings**. In the welcome window, click **Next**. The Host Identification window should be populated automatically. Click **Next**. Click **Finish**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

15. Click **Define deployment options**. In the welcome window, click **Next.** In the Microsoft Update Setup window, click **Use the Microsoft Update service to check for updates (recommended)**, and click **Next**.
16. In the Forefront TMG Protection Features Settings window, verify that **Activate complementary license and enable NIS** is selected in the Network Inspection System (NIS) License box. Verify that the Web Protection License is set to **Activate evaluation license** and that **enable Web Protection** is selected. Verify that **Enable Malware Inspection** is checked, and click **Next.**
17. In the NIS Signature Update Settings window, click **Next.** In the Customer Feedback window, click **No, I don't want to participate**, and click **Next**. In the Microsoft Telemetry Reporting Service window, click the **None** option button. No information is sent to Microsoft. Click **Next**, and then click **Finish**. In the Getting Started Wizard window, uncheck **Run the Web Access wizard**, and click **Close**.
18. Leave your systems running for future projects.

Hands-On Project 10-2: Installing Apache Web Server

Time Required: 10 minutes

Objective: Install Apache.

Description: The open-source Apache program is the most popular Web server in the world. In fact, the most popular current Web server installation is the Apache program running on a Linux operating system. In this project, you install Apache in the Linux system you created in an earlier project.

To install Apache, the Linux system must be connected to the Internet.

1. Start and log on to the Linux system you installed in Hands-On Project 8-1.
2. Click the **Dash home** icon on the left side of the desktop.
3. In the search box, type **terminal**, and press **Enter**.
4. At the command prompt, type **sudo apt-get install apache2**, and press **Enter**.
5. Enter the administrator's password at the prompt.
6. If you are asked whether you want to continue, type **Y**, and press **Enter**.
7. When the installation is complete, click the **Firefox Web Browser** icon on the left side of the desktop, type **http://127.0.0.1** in the URL address box, and press **Enter**. If the installation was successful, you should see a Web page that reads "It works! This is the default web page for this server. The web server software is running but no content has been added, yet."

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

8. Click the **System Settings** icon on the left side of the desktop. Click the **Network** icon. Click **Configure**, and click the **IPv4 Settings** tab. Set the Method list box to **Manual**. Click **Add**, and then configure the IP address as **192.168.1.100**, the netmask as **255.255.255.0**, and the default gateway as **192.168.1.110**. Click **Save**. In the upper-right part of the Wired window, toggle the slider bar to off and then back to on. This action reinitializes the network interface card and displays the new IP address. Close all windows.
9. Leave your system on for the next activity.

Hands-On Project 10-3: Configuring a Web Access Policy

Time Required: 20 minutes

Objective: Configure the basic proxy server functions of TMG.

Description: After installation, TMG blocks all traffic by default. To allow internal clients to access the Internet, a Web access policy modification is required. (In these projects, the Linux Web server functions as the Internet.) In this project, you perform the initial configurations so that the internal Web client, Windows 7, can access the Linux Web server.

1. Verify that the four computers are configured as shown in Table 10-3. Figure 10-20 demonstrates the topology of the lab network.

Computer	IP address	Subnet mask	Default gateway	DNS
Windows 7	10.0.0.110	255.0.0.0	10.0.0.111	10.0.0.125
Windows Server 2008	10.0.0.125	255.0.0.0	10.0.0.111	10.0.0.125
TMG				
Inside	10.0.0.111	255.0.0.0	None	10.0.0.125
Outside	192.168.1.110	255.255.255.0	None	10.0.0.125
Linux	192.168.1.100	255.255.255.0	192.168.1.110	None

© Cengage Learning 2014

Table 10-3 Network configuration

2. Test Web access by opening a Web browser in both the Windows Server 2008 domain controller and the Windows 7 system and then attempting to access http://192.168.1.100. Both attempts should be unsuccessful because the systems are on a different IP segment and the TMG proxy server/firewall is not configured to pass traffic. Repeat this test from the TMG system. This attempt should be successful because TMG has an interface on the same IP segment as the Linux Web server.
3. Verify connectivity by pinging Windows 7 from Windows Server 2008 and by pinging TMG's inside address from both systems. These pings should be successful.
4. In TMG, click **Start**, click **All Programs**, click **Microsoft Forefront TMG**, and click **Forefront TMG Management**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

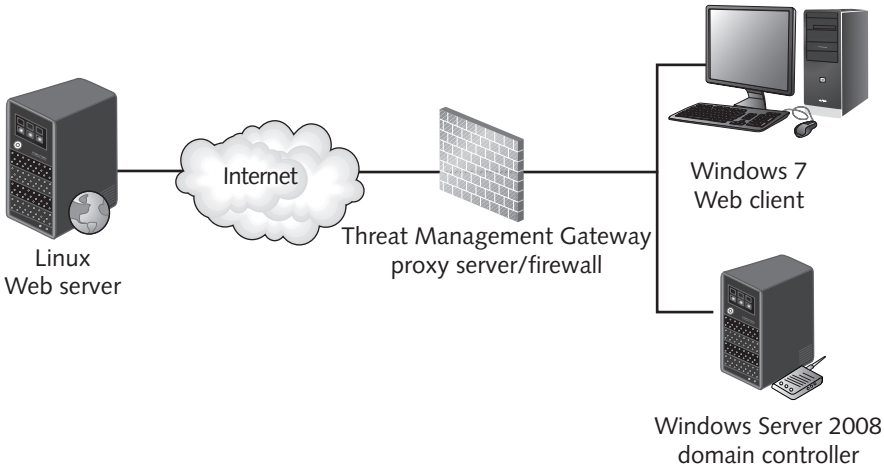

© Cengage Learning 2014

Figure 10-20 Network topology

5. In the Forefront TMG management window, click the plus sign beside Forefront TMG (Team*x*TMG). Click **Web Access Policy**. Notice that the default Web access policy is to block all access to Web sites for internal clients. See Figure 10-21.

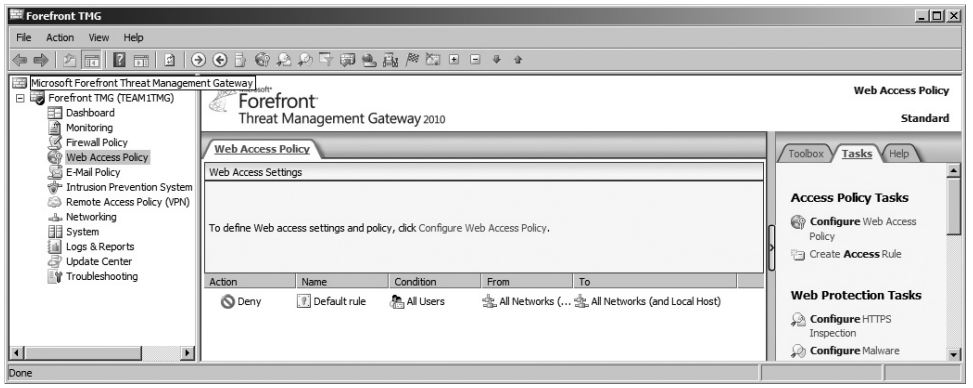

Used with permission of Microsoft Corporation

Figure 10-21 Forefront TMG default Web access policy

6. Click **Configure Web Access Policy**. Read the information in the welcome window, and then click **Next**.
7. In the Web Access Policy Rules window, verify that **Yes, create a rule blocking the minimum recommended URL categories** is selected, and click **Next**.
8. Double-click each of the entries in the "Block access to these Web destinations" box to see a definition of the types of Web sites you will block. Then click **Add**. Expand **URL Categories**, select **Media Sharing**, click **Add**, and click **Close**. Click **Next**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

9. In the Blocked Web Destinations Exceptions window, you can specify users who would not be blocked from accessing the sites you blocked in the previous window. Click **Next**.
10. In the Malware Inspection Settings window, verify that **Yes, inspect Web content requested from the Internet** is selected. Verify that **Block encrypted archives (for example, .zip files)** is checked, and click **Next**.
11. In the HTTPS Inspection Settings window, verify that **Allow users to establish HTTPS connections to Web sites** is selected. Click **Do not inspect HTTPS traffic and do not validate HTTPS site certificates. Allow all HTTPS traffic**. Click **Next**.
12. In the Web Cache Configuration window, uncheck **Enable the default Web caching rule**, and click **Next**. Click **Finish**.
13. Click the **Apply** button. In the Change description text box, type **Baseline Web access policy**, click **Apply**, and then click **OK**. Under Web Access Policy Group, you can see the three rules that are now in effect. These rules are applied in order from top to bottom to determine if a client's Web destination is denied or allowed.
14. From Windows 7, open a Web browser and attempt to access http://192.168.1.100. This attempt should work because the Linux Web server does not match the content categories blocked by the first TMG Web access rule. The second rule, which does apply to this traffic, allows communication.
15. Leave the systems running for the next project.

Hands-On Project 10-4: Examining Logging and Other Features of TMG

Time Required: 15 minutes

Objective: View log file entries and malware inspection functions.

Description: In this project, you examine malware inspection functions and log entries in TMG. For full functionality, note that a subscription is needed to keep intrusion detection, malware, and HTTPS definitions up to date.

1. In TMG, click **Firewall Policy** in the left frame. You should see the same three rules that you created as part of the Web access policy in the previous project.
2. Click **Logs & Reports** in the left frame. Click **Edit Filter** in the right pane. Select **Log Record Type**. In the Value list box, select **Web Proxy Filter**, and then click **Update**.
3. In the Filter by list box, select **Client IP**. Select **Equals** from the Condition list box. In the Value text box, type the IP address of the Windows 7 system. Click **Add to List**.
4. Click **Start Query**. The Logging tab in the center frame shows the progress.
5. From the Windows 7 system, open a Web browser and connect to the Linux Web server, as you did in Step 15 of the previous hands-on project.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

6. In the center frame of TMG, you will soon see the log entry for the allowed connection. Your result should be similar to Figure 10-22. Note that you can determine the rule used to allow traffic, the source and destination IP addresses and ports, and other information.

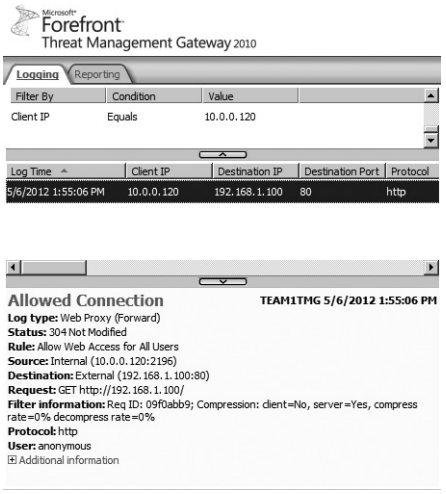

Used with permission of Microsoft Corporation

Figure 10-22 Allowed connection log entry

7. Click **Web Access Policy** in the left frame. Click **Configure Malware Inspection**. Note that the Enable malware inspection check box is selected. Select **Block traffic in relevant rules**.
8. Click the **Content Delivery** tab. Note that the Standard trickling option is selected by default. This option is designed to prevent connection timeouts to the client while traffic is being scanned. As small amounts of traffic are cleared by scanning, they are transferred to the client while the rest of the content is scanned.
9. Click the **Content Types for Progress Notification** button, and then click **Content Types** to view the different types of Web traffic you can configure for progress notification. Click **Cancel**. Click the **Definition Updates** tab, and examine the default settings. This feature is not implemented in this project because a subscription is required.
10. Click the **Inspection Settings** tab. Examine the available options.
11. Click the **Destination Exceptions** tab, and double-click **Sites Exempt from Malware Inspection**. Here you can add domains that can communicate with internal clients without the overhead of malware inspection. Click **Cancel**, and then click **Cancel** again to close the Malware Inspection window.
12. Keep your systems running for the next project.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Project 10-5: Denying Access to Specific Sites for Specific Users

Time Required: 15 minutes

Objective: Configure rule parameters to limit user access to Web functions.

Description: Although Internet access has become a necessity for many organizations, this access also presents problems. Employees might abuse their Internet access and hurt productivity by using social networking sites, or they could even endanger the organization by performing illegal acts through Internet access. TMG allows very specific restrictions to be configured. In this project, you create detailed restrictions for a single user.

1. On the Windows Server 2008 domain controller, create a user with the following attributes:

First name:	Captain
Last name:	Jack
User logon name:	cjack
Password:	Pa$$word

2. In TMG, click **Web Access Policy** in the left frame. Click **Create Access Rule** in the right frame.
3. In the welcome window, type **Web Abusers** in the Access rule name text box, and click **Next**.
4. In the Rule Action window, verify that **Deny** is selected, and click **Next**.
5. In the Protocols window, HTTP and HTTPS are already selected. Click **Add**, expand **All Protocols**, select **FTP**, and click **Add**. Click **Close**, and click **Next**.
6. In the Access Rule Sources window, click **Add**, expand **Networks**, select **Internal**, and click **Add**. Click **Close**, and click **Next**.
7. In the Access Rule Destinations window, click **Add**, click the **New** menu, and click **Domain Name Set**. In the Name text box, type **Social Networking Sites**. Click **Add**, rename the New Domain as **facebook.com**, and click **OK**. From the Domain Name Sets list, select **Social Networking Sites**, and click **Add**. From the **New** menu, click **Address Range**. In the Name text box, type **Linux Web Server**. In both the Start Address text box and the End Address text box, type the IP address of the Linux Web server. Click **OK**. In the Address Ranges list, select **Linux Web Server**, and click **Add**. Click **Close**. Click **Next**.
8. In the User Sets window, select **All Users**, and click **Remove**. Click **Add**, and click the **New** menu. In the welcome window, type **Web Abusers** in the User set name text box, and click **Next**. In the Users window, click **Add**, and click **Windows users and groups**. In the Select Users or Groups window, click the **Locations** button. If the Windows Security window appears, enter **administrator** as the username and **Pa$$word** as the

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

password, and click **OK**. In the Locations window, expand **Entire Directory**, select **Team*x*.net**, and click **OK**. In the Enter the object names to select text box, type **cjack**, and click **Check Names**. Captain Jack's name should appear underlined. Click **OK**, click **Next**, and click **Finish**. In the Add Users window, select **Web Abusers**, click **Add**, click **Close**, and click **Next**. Click **Finish**. Click **Apply**. In the Change description text box, type **Web Abusers created and enabled**. Click **Apply**, and click **OK**.

9. From the Windows 7 system, log on to the domain as Captain Jack. Open a Web browser and attempt to access the Linux Web server, as you did in Step 15 of Hands-On Project 10-3. Your attempt should fail. If your systems were connected to the Internet, access to Facebook and FTP sites would also be denied.
10. Log off the Windows 7 system, and log on as another domain user. Attempt to access the Linux Web server again. This attempt should succeed.
11. Log off all computers.

Case Projects

Case Project 10-1: Firewall Selection Parameters

You work for a network consulting firm and you have been asked to create a guide for firewall selection. Create an outline for a section that lists and defines the most important parameters of firewall functions, such as throughput.

Case Project 10-2: Proxy Server Parameters

You are still working on your guide for firewall selection from the previous case project. Create an outline for a section that lists and defines the most important functional parameters of proxy servers.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 11

VPN Concepts

After reading this chapter and completing the exercises, you will be able to:

- Explain basic VPN concepts
- Describe encapsulation in VPNs
- Describe encryption in VPNs
- Describe authentication in VPNs
- Explain considerations for planning VPN deployments
- Describe options for VPN configurations
- Explain ways to adjust packet-filtering rules for VPNs
- Describe ways to review VPN policies and procedures

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

A virtual private network (VPN) combines two essential elements of a network defense strategy: security and connectivity. VPNs play important roles in businesses that rely on the Internet for critical communications and are attractive because of their low cost and efficiency.

In the past, companies that conducted electronic transactions used leased lines—private connections rented from telecommunications companies, using technologies such as frame relay or T-carriers (T1 and T3). Leased lines are expensive, however, so a VPN is a cost-effective way for networks to make a secure connection through the public Internet.

A VPN is said to be "virtual" because the connection between networks does not use a dedicated line. Rather, it uses the same public Internet connections that millions of people use. Users who make VPN connections gain privacy via a variety of technologies, such as encryption and authentication. In this chapter, you learn what VPNs are, why they are growing in popularity, and how they ensure private communication between networks. This chapter also explains the protocols that create a secure virtual tunnel between computers and describes encryption methods for maintaining the privacy and integrity of VPN communications. You will learn how VPNs are deployed, how they are integrated with firewall configurations, and about relevant policies and procedures.

Understanding VPN Concepts

A **virtual private network** (**VPN**) provides a way for two computers or computer networks to communicate securely by using the same public communication channels available on the Internet, where millions of computers and networks exchange data.

The Internet is the most common network used for VPNs, but you can create a VPN on any network, large or small.

To understand what VPNs are, consider how regular mail works. Because you want your letter kept private, you place it in a sealed envelope (encapsulation). If you are concerned about the possibility of your letter being intercepted and read by someone else, you could write it in code (encryption), but you would need a way to let the receiver decrypt it, which could be considered a "key exchange."

You need to tell the mail carrier where to deliver the envelope and who should receive it, so you include a name and address (the destination address information). In case your envelope cannot be delivered, you put a return address on it (the source address) so that the mail carrier can bring it back to you. As on the public Internet, your message can take several paths to its destination. This is the "virtual" part of a VPN.

Finally, the mail is processed through many different sorting centers, air and ground transportation systems, and post offices in the postal service "network," just as the Internet is a conglomeration of private LANs, public transmission lines, and other systems that form a giant mesh network.

Specified computers, users, or network gateways are identified as endpoints of the VPN connection, which is called a **tunnel**; only those designated computers, users, or gateways can participate

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

in the VPN (see Figure 11-1). A VPN is a network because it connects computers and extends an organization's network beyond its current boundaries. A VPN, then, is a virtual network that uses the Internet to establish a secure connection. VPNs enable computers to exchange private encrypted messages that others cannot decipher.

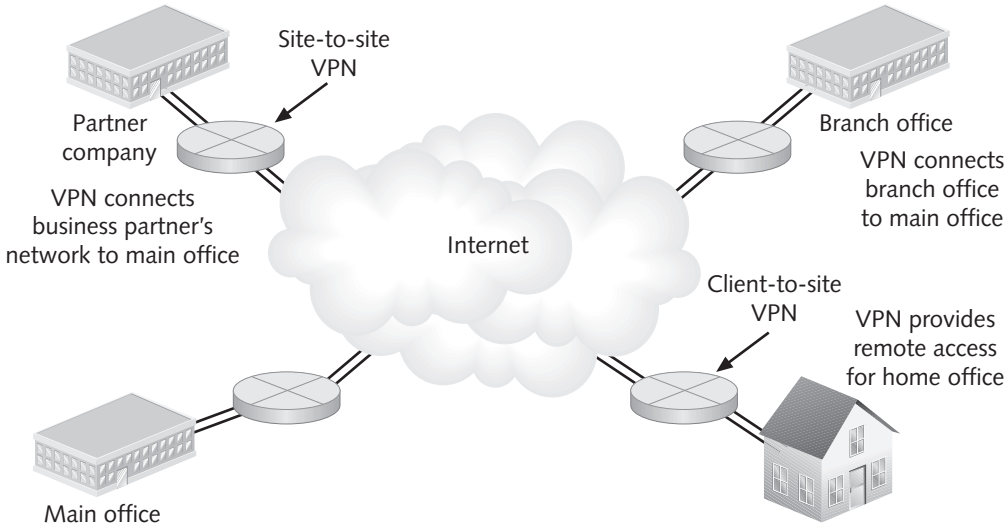

© Cengage Learning 2014

Figure 11-1 Establishing connections with a VPN

VPN endpoints represent extensions of participating networks. If these endpoints are not secured by a firewall, they could give intruders a way to access the network. Unless your VPN client software incorporates its own firewall, you need to make sure any remote computers that connect to your organization's VPN are equipped with desktop firewalls.

VPN Components

VPNs can be assembled using a variety of components. However, all VPNs contain some essential elements that enable data to be transmitted securely from one point to another:

- *VPN server or host*—A **VPN server** is configured to accept connections from clients who connect via dial-up or broadband.
- *VPN client or guest*—A **VPN client** can be a router that serves as the endpoint of a site-to-site VPN connection, which uses hardware to connect two networks. It can also be an operating system (OS) configured to function as an endpoint in a VPN.
- *Tunnel*—The connection through which data is sent.
- *VPN protocols*—**VPN protocols** are groups of standardized communication settings that software and hardware use to encrypt data sent along the VPN. They include Internet Protocol Security (IPsec), Point-to-Point Tunneling Protocol (PPTP), Layer 2 Tunneling Protocol (L2TP), and Secure Sockets Tunneling Protocol (SSTP).

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The number of components in a VPN depends on the number of networks in its configuration. For instance, if a VPN contains four networks, it has at least four separate servers and four tunnels; clients from each endpoint can participate in any of the VPN tunnels they have permission and credentials to access.

In general, you can set up two different types of VPNs. The first type links two or more networks and is called a **site-to-site VPN** (or a gateway-to-gateway VPN). The second type makes a network accessible to users who need remote access, and is called a **client-to-site VPN** (or a remote access VPN).

Some companies that maintain VPNs with partner organizations benefit by using the same ISP as their partners for an Internet connection. Positioning participants in the VPN on the same part of the Internet backbone can make the VPN run more smoothly and reliably.

Types of VPNs

Hardware VPNs The components you choose to establish a VPN depend on whether you want to use existing hardware or software. Creating a VPN with new components increases costs but has the benefit of reducing the load on other network security components, such as firewalls.

Hardware-based VPNs connect one gateway to another to create a gateway-to-gateway VPN. Typically, the VPN hardware is a router at each network gateway that encrypts outbound packets and decrypts inbound packets. Another hardware option involves a **VPN appliance**, a hardware device designed to serve as the VPN endpoint and join LANs (see Figure 11-2).

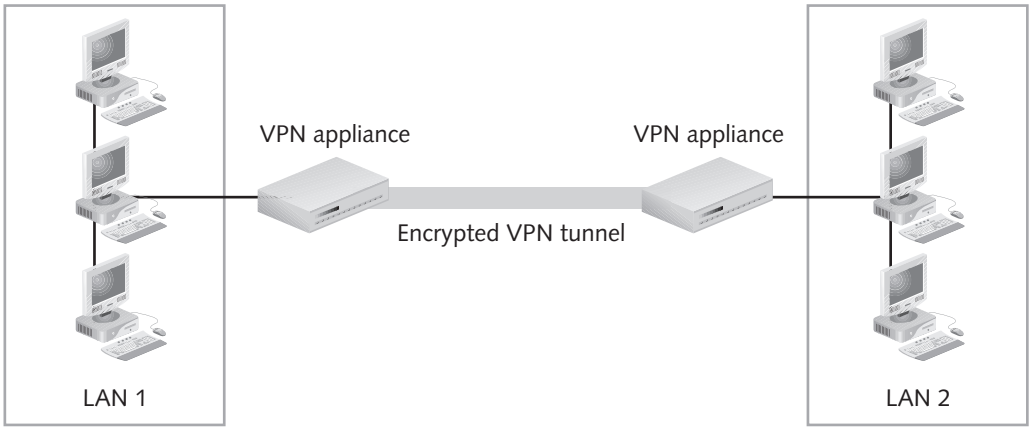

© Cengage Learning 2014

Figure 11-2 VPN appliance creates secure connections between two or more LANs

In general, hardware products such as routers that are dedicated to VPN operation tend to handle more network traffic than software products. As a result, they are more scalable than software VPNs. They can also be more secure than software VPNs because they are designed from the ground up for a specific purpose, and they do not depend on an underlying OS

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

that might have security flaws. Hardware VPNs should be the first choice for fast-growing networks that need to encrypt all traffic passing through the VPN device. They are also a good choice when VPN endpoints use the same type of routers and are controlled by the same organizations.

The security industry has moved toward integrated solutions, as discussed previously. The most common term you see with integrated security products is **Unified Threat Management (UTM)**. UTM describes how to create a single product suite by integrating many security-related functions, including firewalls, intrusion detection and prevention, antivirus, e-mail, and communication security, along with VPNs, secure remote access, and authentication. Many excellent products provide these functions in a single platform. Integrated solutions have some major advantages, particularly cost, interoperability, and easier management. On the other hand, relying on a single device introduces a single point of failure and reduces layers of defense. In reality, even though a UTM product can perform multilayered defense functions, it is still a single product. An organization's security needs dictate whether a UTM product is appropriate.

Software VPNs Most software-based VPNs are integrated with firewalls and are more cost-effective than hardware VPN devices. Software-based VPNs also increase network security because they are integrated with functions that a firewall already performs, such as packet filtering and Network Address Translation (NAT). Software-based VPNs are appropriate when participating networks use different routers and firewalls or when the endpoints are controlled by different organizations and network administrators.

The main reason for using software VPNs to link networks is the flexibility they offer. They can be configured to enable traffic based on domain name, protocol, or IP address. These restrictions prove useful when some, but not all, of the traffic passing through the VPN is meant to be encrypted and sent through the tunnel. However, because software VPNs often rely on the OS on which they are installed, configuring and using them are complex tasks.

Evaluating Business Needs for VPNs

Planning a VPN deployment requires assessing an organization's goals. For example, the need to keep business transactions private has prompted more organizations to adopt VPNs. The popularity of e-commerce is an incentive as well, and government and military agencies need to share information to provide effective homeland security.

A VPN is an excellent solution for an organization that needs to follow a budget while maintaining security. Budgetary considerations have always made VPNs an attractive business proposition. When you use a VPN, you are essentially spreading the cost of its operation over many users, which makes it cost effective.

In addition, many companies employ remote contractors who need to access the corporate network from their homes or offices. Employees who travel for business reasons need to check e-mail and exchange files with colleagues in the central office. Secure remote access is an essential requirement for many businesses and is an important reason for establishing a VPN. Another reason is the need for a secure means of connection for partners, suppliers, contractors, and others outside the company who need real-time data access to support just-in-time processes, inventory management, and shipment status information.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

You should review the company security policy for guidance on existing security goals and procedures, and as a basis for planning VPN deployment. Also, you should consider the type of business an organization runs, along with its number of employees and existing infrastructure, as well as security and throughput rates required for data. With this information, you can integrate a VPN with minimal disruptions and avoid unexpected and undesirable outcomes.

When considering the type of business, ask yourself what the business does. What product or service does it sell, and who are its customers? To set up an effective VPN, you must know who needs access and where VPN users are located.

Analyzing the existing infrastructure helps you determine how and where to place a VPN. If the network administrator has an accurate network map, this task is fairly easy. Often this documentation is unavailable, however, so you might need to examine the network and make your own map. It does not need to be complex, but it must accurately reflect the environment into which you are integrating VPN services.

After deciding on the placement of a VPN, you can begin researching hardware or software. Determine whether you should reconfigure existing hardware to support VPN connections or purchase a VPN appliance, server, firewall, or router to do the job. Your assessment of the organization's needs and existing infrastructure can guide you in this decision, and vendors are always happy to assist potential buyers. Use their knowledge to help you decide what tools will meet your needs, offer sufficient scalability, stay within budgetary constraints, and integrate with existing systems. You might want to develop a list of requirements when meeting with vendors to ensure that nothing is overlooked. Remember to follow security policy guidelines, however, when giving out network information. Some information, such as trade secrets or business processes that give your company a competitive advantage, must remain confidential. For this reason, most companies require a nondisclosure agreement. This document is a legally binding contract that stipulates what information the parties will share, and it requires the vendor to maintain confidentiality.

Cost is also a key factor, and it often narrows the choices of hardware and software. Beware of "bleeding-edge" technology that could become outdated or unsupported just months after you purchase it. Also, realize that having unlimited funds does not mean your design will be foolproof. Regardless of budgetary factors, your design must integrate with existing technology and provide services without failure. In other words, you do not need to reinvent the wheel. For example, if you are installing VPN services on a network with older firewalls on the perimeter, you could configure those firewalls to provide VPN services instead of replacing them with newer ones. If the firewalls provide the necessary services with acceptable performance and security and offer adequate scalability, upgrading probably is not necessary.

VPNs are as diverse as the needs they support. When you have a clear picture of the organization and its users, you can begin planning the VPN's configuration, testing, and deployment. A secure VPN design should address the following factors:

- Secure connectivity
- Encryption
- Availability
- Authentication

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Secure management
- Reliability
- Scalability
- Performance

Advantages and Disadvantages of VPNs

VPNs have advantages and disadvantages. They offer a high level of security, provided that network administrators address inherent challenges. For instance, if a VPN device is configured incorrectly or remote users at a VPN endpoint disable their firewalls by mistake and let in an attacker, the VPN's protections can be circumvented. In addition, VPNs can be complex to configure, and the necessary hardware and software can represent a substantial investment. Table 11-1 summarizes some of the main advantages and disadvantages.

Advantages	Disadvantages
Far less expensive than leased lines	Can be complex to configure
Many elements working together provide strong security	Can result in slower data transfer rates than a leased line
Standards and protocols used in VPNs are well developed and widely used	Depends on the often unpredictable Internet; if your ISP or other parts of the Internet go down, your VPN goes down
Can result in less overall complexity in an organization's network	Requires administrators to install VPN client software on remote computers
Can make use of a company's existing broadband connection	VPN hardware and software from different vendors might prove incompatible because they use different protocols

© Cengage Learning 2014

Table 11-1 Advantages and disadvantages of VPNs

By focusing on Internet-based technologies, VPNs simplify a network. You have only one Internet connection to manage instead of managing an Internet connection plus leased lines. In addition, running a VPN means you have even more ways to maximize network uptime. Downtime is expensive; in addition to the cost of repairs, it adds administrative time in troubleshooting and repairing problems and affects users' productivity.

The Three VPN Core Activities

Encapsulation

VPNs can use unsecure, public Internet connections and still provide a high level of security because they perform a core set of activities: encapsulation, encryption, and authentication. Together, these activities tunnel data from one network to another using the infrastructure of the Internet.

First, VPNs perform **encapsulation** of data: They enclose a packet within another packet that has different IP source and destination information for a high degree of protection.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Encapsulation protects the integrity of data sent through the VPN by hiding the data packets' source and destination information. The VPN encapsulates data packets in packets that use the VPN gateway's source and destination addresses, as shown in Figure 11-3. The gateway could be a router that uses IPsec, a VPN appliance, or a firewall that functions as a VPN and has a gateway setup.

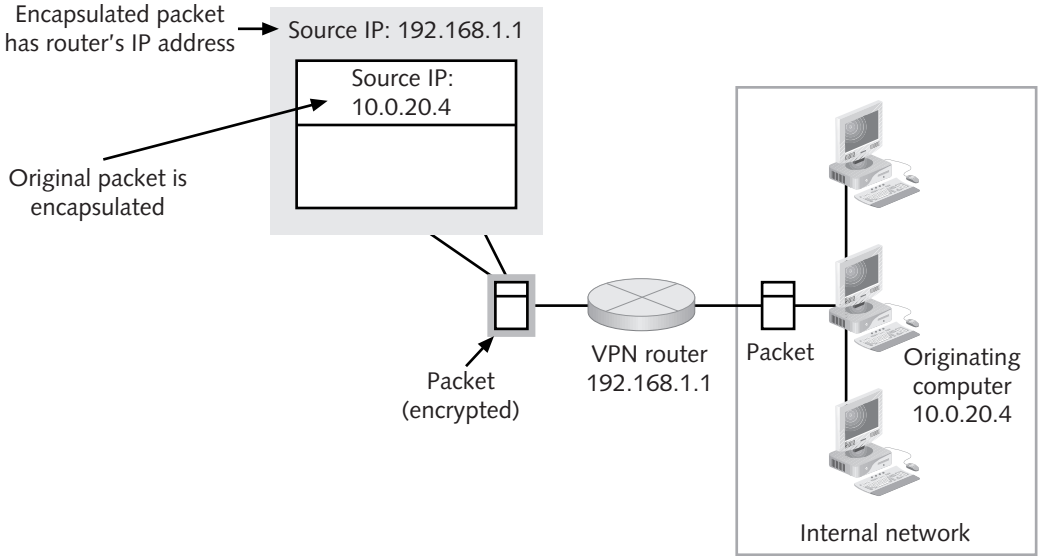

© Cengage Learning 2014

Figure 11-3 Encapsulating data to conceal source and destination information

When a VPN tunnel is in place, the source and destination IP addresses of the encapsulated data packets can be in private reserved blocks that are not routable over the Internet, such as the 10.0.0.0/8 addresses or the 192.168.0.0/16 reserved network blocks.

Understanding Tunneling Protocols

Because VPNs link networks and gateways that might have different OSs or hardware platforms, protocols need to be established so that communication can take place. When you configure a device to function as a VPN host, you need to choose the protocol you want to use. For instance, IPsec with Internet Key Exchange (IPsec/IKE) has become a very common method of establishing a VPN. Also, Secure Sockets Layer can leverage Web-based applications and is becoming very popular. A number of common VPN protocols are discussed in the following sections.

Point-to-Point Tunneling Protocol

IPsec provides VPN security in many circumstances, but it is not appropriate for every application. Users who need to dial up a server on computers with older OSs might need to connect to VPN servers configured to support **Point-to-Point Tunneling Protocol (PPTP)**.

PPTP encapsulates TCP/IP packets and uses a proprietary technology called Microsoft Point-to-Point Encryption (MPPE) to encrypt data passing between the remote computer and the remote access server. PPTP should be used only in trusted networks where security is addressed

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

by other means. Vulnerabilities in MPPE and Microsoft Challenge/Response Authentication Protocol (MS-CHAP) make PPTP a poor choice for high-performance networks with many hosts. Use PPTP only for a small-scale VPN that needs to support mobile users. The best option is to upgrade older clients to use more secure protocols when possible.

Layer 2 Tunneling Protocol **Layer 2 Tunneling Protocol** (**L2TP**) is based on two older PPP tunneling protocols, Cisco's Layer 2 Forwarding (L2F) and PPTP, and it uses UDP port 1701. Using L2TP, a host machine can connect to a modem and have its PPP data packets forwarded to another remote access server. L2TP encapsulates the PPP packets. When the data reaches the remote access server, the payload is unpacked and forwarded to the destination host on the internal network.

The endpoints of an L2TP tunnel are called the L2TP Access Concentrator (LAC) and the L2TP Network Server (LNS). Generally, the LAC is the initiator of the tunnel, and the LNS is the server waiting for tunnels (although either side can initiate tunnels). After an L2TP tunnel is established, it is bidirectional. L2TP exchanges two types of packets within the tunnel: control and data packets. Reliability features are provided for control packets but not for data packets. Reliability must be provided for data packets by higher-layer tunneled protocols.

When used alone, L2TP does not provide confidentiality or authentication, so it is usually combined with IPsec for improved security. This combination is referred to as L2TP/IPsec, and is a more secure alternative to PPTP. IPsec enables L2TP to perform authentication and encapsulation as well as encryption. It encapsulates the entire L2TP packet, which has a couple of notable benefits. Because the L2TP packet is encapsulated in the IPsec packet, no information about the internal network can be gained from the encrypted packet. Also, the hidden L2TP packet is contained in the IPsec packet until it reaches the endpoint, so UDP port 1701 does not need to be opened unless a firewall or packet filter is acting as an endpoint. (IPsec is explained in detail later in this chapter in the "Understanding IPsec/IKE" section.)

Generic Routing Encapsulation **Generic Routing Encapsulation** (**GRE**) is a tunneling protocol that can encapsulate a variety of Network layer packets inside IP packets. GRE encapsulates the actual packet inside a GRE packet, which is then encapsulated in an arbitrary transport protocol for delivery. In IPv4 networks, the transport protocols are TCP and UDP, of course. Other protocol stacks can also work with GRE, which is why it is called "generic."

GRE is a stateless protocol, meaning that endpoints do not monitor the status of other endpoints. This is useful because an endpoint might not be able to detect a service provider or network's internal tunneling architecture. Cisco designed GRE to create virtual point-to-point links to Cisco routers at remote locations, but it can be used for several applications. For example, it can be combined with PPTP and IPsec and is commonly used with PPTP in Microsoft VPNs.

Here is how GRE works:

1. The packet to be transmitted is passed to the local endpoint of the VPN tunnel.
2. The packet is given a PPP header and encrypted.
3. GRE encapsulates the packet.
4. A transport protocol header is added.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

5. The packet is transmitted through the tunnel.
6. At the endpoint, the transport protocol and the GRE and PPP headers are stripped off.
7. The payload may or may not be decrypted, depending on the endpoint configuration and encryption used.
8. The payload is transmitted to its destination.

GRE uses IP protocol 47 and native IPv4 routing. Route filtering is not changed, but if packet filters are used, the filter must be configured to examine inside the GRE packet or filtering should be done at tunnel endpoints. In such cases, it might be wiser to terminate the tunnel at the firewall. IPv6 is not currently included as a delivery or payload.

GRE can enable attacks against VPN client computers. The GRE protocol is not attacked, but attackers can build a GRE packet and use it to initiate a connection to the client. Attacks on a protocol level require some effort to discover basic information about the target network. When used correctly, this type of attack can enable an intruder to gain a legitimate connection. For more information on attacks that use GRE, visit *www.phenoelit-us.org/irpas/gre.html* or conduct a Web search for GRE protocol attacks.

Understanding IPsec/IKE

Internet Protocol Security (IPsec) is a set of standard procedures that the Internet Engineering Task Force (IETF) developed for enabling secure communication on the Internet. IPsec has become the standard set of protocols for VPN security for a number of reasons:

- IPsec works at Layer 3 and provides a type of security that is not available with protocols that work at Layer 2.
- IPsec can encrypt an entire TCP/IP packet instead of just the data portion, as with other protocols.
- IPsec was originally developed for use with IPv6, although it can also work with IPv4.
- IPsec provides authentication of source and destination computers before data is encrypted or transmitted.

Perhaps the biggest advantage of using IPsec is that it has been standardized and is supported by a wide variety of VPN hardware and software devices. OSs such as Windows Server 2008 and Windows 7 enable you to set up an IPsec connection with another Windows computer that has IPsec enabled.

When an IPsec connection is established between two computers, they authenticate one another and then establish the **Security Association (SA)** settings they use to communicate. An SA is a relationship between two or more entities that describes how they will use security services to communicate. Each IPsec connection can perform encryption, encapsulation, authentication, or a combination of the three. When determining which services to use, the entities in a connection must agree on the details, such as which algorithm to use for encryption. After the transaction is completed, the entities must share session keys. An SA is the method that IPsec uses to track all the details of a communication session. An SA is

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

unidirectional, meaning that it is set up in each direction of a communication, forming two one-way SAs between entities.

These transactions take place in the background. However, in an OS environment, you need to decide whether IPsec is required for all connections to the host machine, or whether the host requests an IPsec connection for computers or other devices that support it. If IPsec is not supported on the client machine, it is not used. You also have the option of requiring a secure connection to another computer over a VPN. If you want to connect to another computer while requiring IPsec, you need to adjust the packet-filtering rules; otherwise, your IPsec-enabled computer blocks all other connections by default.

IPsec Components IPsec's many components provide encryption, encapsulation, key exchange, and authentication. These components include the following:

- **Internet Security Association Key Management Protocol (ISAKMP)** enables two computers to agree on security settings and establish an SA so that they can exchange keys by using Internet Key Exchange.
- The **Internet Key Exchange (IKE)** protocol enables computers to exchange keys to make an SA. By default, IKE uses UDP port 500 on both the client and server. Different configurations might use different ports. For example, IPsec NAT traversal uses UDP ports 500 and 4500 for IPsec traffic. You will learn more about NAT traversal later in this chapter.
- The Oakley protocol enables IPsec to use the Diffie-Hellman encryption algorithm to create keys. (You can learn more about Diffie-Hellman encryption at *www.ietf.org/rfc/rfc2631.txt*.)
- The IPsecurity Policy Management service runs on Windows computers. It retrieves IPsec security policy settings from Active Directory and applies them to computers in the domains that use IPsec.
- An **IPsec driver** is software that handles the tasks of encrypting, authenticating, decrypting, and checking packets.

Suppose you have configured a VPN connection between two computers and you want it to use IPsec. When one IPsec-compliant computer connects to the other, the following events occur:

1. The IPsec driver and ISAKMP retrieve the IPsec policy settings.
2. ISAKMP negotiates between hosts, based on their policy settings, and builds an SA between them.
3. The Oakley protocol generates a master key to secure IPsec communication.
4. Based on the security policy established for the session, the IPsec driver monitors, filters, and secures network traffic.

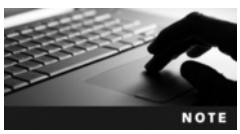

IPsec is not foolproof. For instance, if the machine that is running IPsec-compliant software has already been compromised, no communication from it, including IPsec communication, can be trusted. IPsec is not a substitute for firewall, antivirus, and intrusion detection software.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

IPsec in tunnel mode is not recommended for remote access VPNs. For this purpose, use L2TP/IPsec.

The two core IPsec components are the ones that protect TCP/IP packets exchanged in the VPN: Authentication Header and Encapsulating Security Payload. Both are discussed in the following sections.

Authentication Header

Authentication Header (**AH**) is an IPsec component that authenticates TCP/IP packets to ensure data integrity. With AH, packets are signed with a digital signature that tells other IPsec-compliant devices the packet contains accurate IP header information because it originated from a computer using IPsec. Digitally signing a packet indicates that it has not been tampered with and that the IP information in the header has not been spoofed. It ensures integrity but not confidentiality.

To authenticate all or part of a packet's contents, AH adds a header that is calculated by IP header and data values in the packet, essentially creating a message digest of the packet. Security is achieved by calculating the values with a hashing algorithm and a key known only to entities in the transaction. (Keys are negotiated and exchanged when the SA is set up.) Because only authorized entities have the key, only they know how to run the algorithm to see whether the data has been tampered with. If the values match, the message's authenticity is considered intact.

An AH header does not change the contents of a message. It simply adds a field following the IP header. The field contains the data and the computed value of the IP header, except any fields that change in transit, such as the time to live (TTL) field, as shown in Figure 11-4.

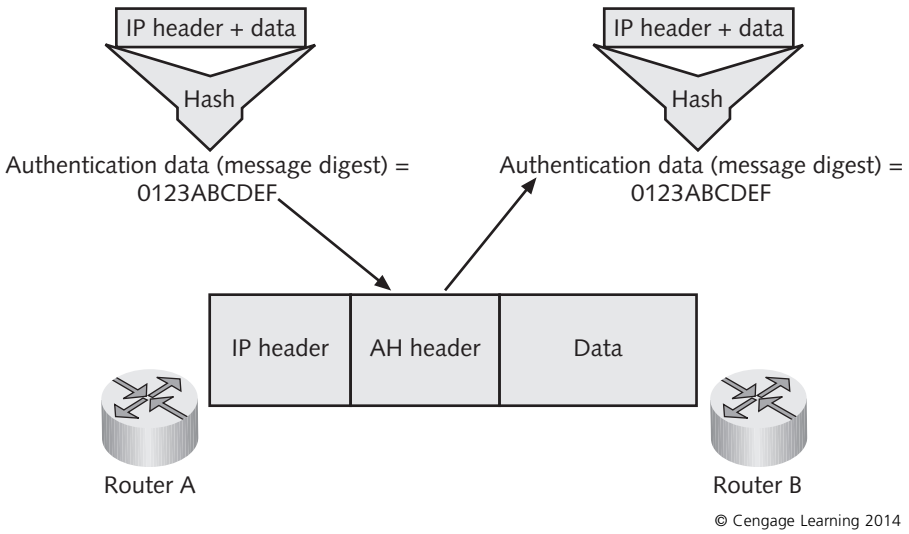

Figure 11-4 AH message exchange

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

AH works a little differently in the two IPsec modes: tunnel and transport (see Figure 11-5). In tunnel mode, AH authenticates the entire original header and builds a new IP header that is placed at the front of the packet. The only fields not authenticated by AH in tunnel mode are fields in the new IP header that can change in transit. In transport mode, AH authenticates the data and the original IP header, except fields that change in transit.

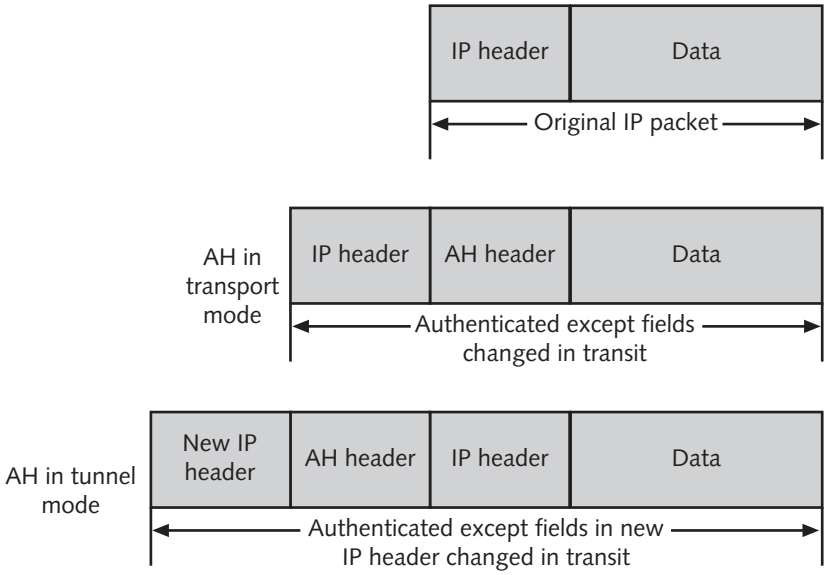

© Cengage Learning 2014

Figure 11-5 AH in tunnel and transport modes

Encapsulating Security Payload As mentioned, AH ensures authentication and integrity for messages but not confidentiality. To ensure the confidentiality of data transmitted through a VPN tunnel using IPsec, you use **Encapsulating Security Payload (ESP)**. ESP encrypts different parts of a TCP/IP packet, depending on whether IPsec is used in transport or tunnel mode.

In tunnel mode, ESP encrypts both the header and data portion of a packet. This encryption protects data, but because the IP header is encrypted as well, the data cannot pass through a firewall that performs NAT; the firewall does not know how to interpret the IP source and destination information in its encrypted form. In transport mode, only the data portion is encrypted. As a result, if the VPN is used with a firewall that performs NAT, IPsec should be configured to work in transport mode. Figure 11-6 shows the difference the IPsec mode makes to ESP.

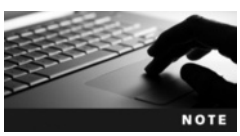

Using AH and ESP together offers additional security, but you might not want to use ESP if another device or application is already providing encryption. ESP, like other forms of encryption, requires substantial processing resources and can slow the rate of data transfer through a VPN.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

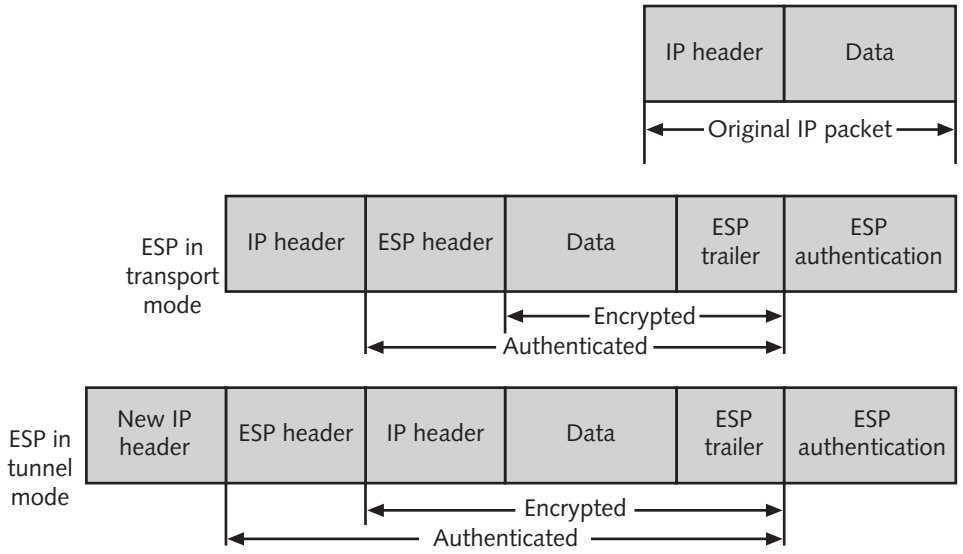

© Cengage Learning 2014

Figure 11-6 ESP in tunnel and transport modes

Multiprotocol Label Switching (MPLS) is an IETF initiative that provides efficient routing, forwarding, switching, and designation of traffic. It is independent of Layer 2 and 3 protocols and provides a way to map IP addressing information into a simple, fixed-length label. After a packet has been labeled, it is routed by label switching. MPLS shows promise for reducing latency and easing strained Internet backbone infrastructures, and does not rely on encapsulation or encryption for security. Although MPLS is beyond the scope of this book, it is worthwhile to learn more about the technology. You can read more about MPLS (RFC 3031) at *www.faqs.org/rfcs/rfc3031.html*.

Encryption

As you learned in Chapter 5, encryption is the process of rendering information unreadable by all but the intended recipient. The encryption process is carried out by means of an algorithm that generates an encoded block of data called a key. The key is part of an electronic document called a digital certificate, which is obtained from a certification authority (CA), a trusted organization that issues keys. The key is then used to encrypt data at the originating endpoint of the VPN and to decrypt it at the destination endpoint (see Figure 11-7).

To perform encryption at both endpoints of the VPN, the keys must be exchanged by participants who have an SA. The exchange can be performed by using a variety of encryption methods. In symmetric cryptography, the same key is exchanged by sender and recipient. In asymmetric cryptography, two different keys are used—a public key and a private key. When a person or an organization obtains a digital certificate from a CA, an encryption algorithm is used to generate a private key. This key is never exchanged but is maintained

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

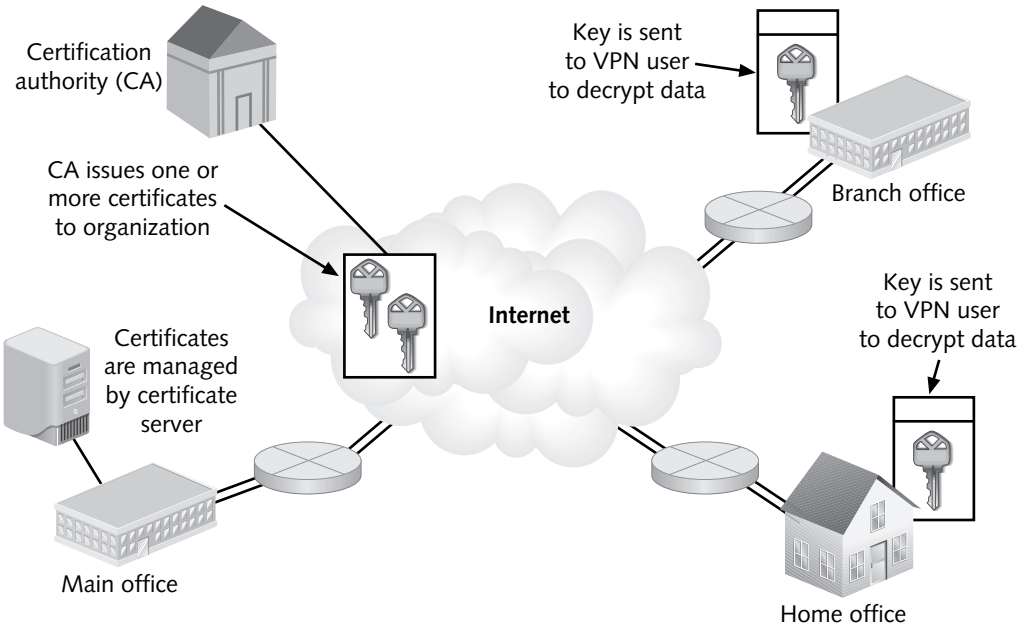

Figure 11-7 VPN endpoints encrypt and decrypt data by exchanging keys

securely by the certificate holder. The private key is used to generate a public key, which can be exchanged freely among VPN participants.

Another key exchange method, IKE, uses **tunnel method encryption** to encrypt the header and data components of a packet and to encapsulate the packet within a new packet that has a different header. IKE is increasingly popular because it provides a high level of security, which outweighs the decrease in network performance caused by complex encryption.

Encryption Schemes One advantage of a VPN is its capability to extend a wide area network (WAN) to multiple locations by using the Internet. However, the openness of the Internet creates security risks that need to be addressed by encryption. Encryption is one of the most important aspects of a firewall because a VPN uses encryption to achieve its primary goal of preventing unauthorized users from reading the data payload.

Encryption schemes do not ensure a uniform level of security, however. Most older encryption algorithms, such as Data Encryption Standard, have been cracked. The reason is simple: As computers have grown faster and more powerful, more resources can be brought to bear on cracking algorithms. The approach is usually brute force, meaning that all possible keys are tried. Generally, the longer the key, the more difficult the algorithm is to crack. Therefore, algorithms had to become more complex, use longer keys, and use more rounds of computations. Some encryption schemes provide strong encryption that uses long keys (for instance, 128-bit keys) or multiple keys. The following sections describe some encryption methods commonly used in VPNs.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Triple Data Encryption Standard Many VPN products use Triple Data Encryption Standard (3DES) encryption. 3DES is a variation on Data Encryption Standard (DES), an older protocol standard IBM developed in the mid-1970s that was adopted as an encryption standard in 1977. DES is not a secure encryption method; it can be cracked if enough computers are working on it. For this reason, it is considered obsolete and is rarely used now. 3DES is far more secure because it uses three separate 64-bit keys to process data. The first key encrypts the data, the second key decrypts it, and the third key encrypts it again (see Figure 11-8).

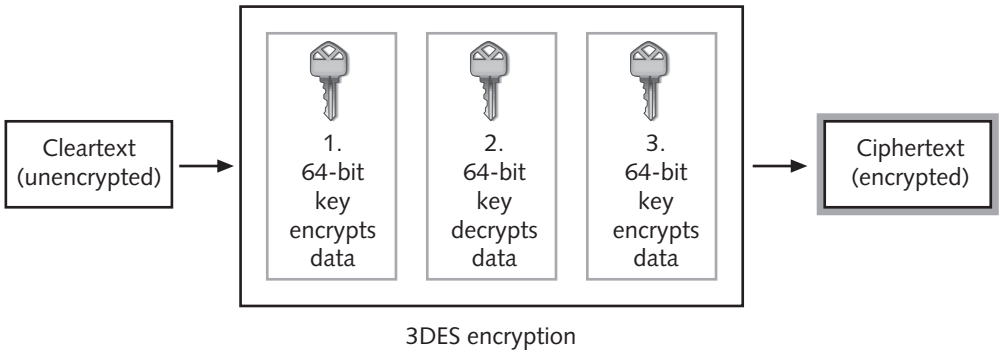

© Cengage Learning 2014

Figure 11-8 Many VPNs use 3DES encryption

Although using three keys results in a strong level of encryption, extensive time and resources are needed to encrypt the information. 3DES encryption takes three times as long as DES, which uses only a single key. 3DES is also more processing intensive, making its use a trade-off. Although most modern computers can use stronger encryption without intolerable performance degradation, encryption must be selected and managed carefully. Use stronger encryption to ensure confidentiality when needed and other forms of security, such as digital certificates, for authentication or integrity. In short, use encryption only when it is necessary.

Advanced Encryption Standard The National Institute of Standards and Technology (NIST) developed a new encryption standard, Advanced Encryption Standard (AES), to replace 3DES as the U.S. government standard in 2002. AES is stronger than 3DES and works faster. It uses the Rijndael symmetric encryption algorithm, which specifies how to use 128-bit, 192-bit, or 256-bit keys on 128-bit, 192-bit, or 256-bit blocks. The standard AES implementation has a fixed block size of 128 bits and uses keys of 128, 192, or 256 bits.

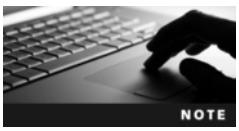

In practice, the Rijndael algorithm can support a larger range of key and block sizes than AES.

AES is a block cipher, meaning that mathematical computations are applied to blocks of data. Each iteration (application of the algorithm) is called a round. AES applies 10 rounds

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

for 128-bit keys, 12 rounds for 192-bit keys, and 14 rounds for 256-bit keys. By comparison, other block ciphers use 7, 8, and 9 rounds, respectively, against the best-known attacks.

One concern about AES centers on its mathematical structure. Most ciphers do not have the neat algebraic description of AES, leading some cryptographers to speculate that a flaw in its mathematical structure could lead to successful attacks. As of this writing, some successful attacks have been launched against certain AES implementations, and others are theoretically possible. One successful type of attack is a **side channel attack**, which does not attack the cipher directly; instead, it attacks the underlying systems that leak data. Some classes of side channel attacks include the following:

- *Timing attacks*—Attacks based on measuring how much time computations take to perform. For example, a timing attack might monitor data moving in and out of the CPU or memory. Monitoring how long it takes to transfer key information can yield clues about key length or eliminate certain key lengths.
- *Power-monitoring attacks*—Attacks that use varying power consumption levels by hardware during computation. Watching the power input to the CPU during computation offers information that can help determine the algorithm.
- *Acoustic cryptanalysis*—Similar to power-monitoring attacks, acoustic cryptanalysis uses the sound produced by computations. The current used to power hardware produces heat, which is leaked into the atmosphere. The fluctuations of heating and cooling (thermodynamics) produce low-level acoustic noise that can be examined for clues about the underlying system.
- *Radiation monitoring (also known as van Eck or TEMPEST)*—Leaked radiation provides plaintext or other information that can be used to launch attacks. Electrical current fluctuations generate radio waves, which can occur in patterns. The patterns can be recorded and analyzed to gain information about associated hardware, and sometimes bits of actual data can be captured.
- *Thermal imaging attack*—If the surface of the CPU can be seen, infrared images can be taken that provide clues about the code being run.

Side channel attacks rely on emitted information, as in acoustic or radiation monitoring, and relational information, as in timing or power-monitoring attacks. Countermeasures against side channel attacks include power conditioning and UPS systems to control power fluctuations and emissions, shielding to prevent radiation leakage, and strong physical security to prevent acoustic recorders or other monitoring devices from being installed.

Another possible attack against AES is an XSL attack, which is extremely complex and uses multivariate quadratic equations. However, the complexity makes the attack unlikely to affect AES in practice in the foreseeable future. You can learn more about XSL attacks (and find more information on cryptography) at *www.cryptosystem.net/aes*.

Launching these attacks requires a high level of expertise, which eliminates many potential attackers. As successful attacks emerge, however, automated scripts and instructions invariably appear, making it possible for less knowledgeable attackers to use the techniques. Keep in mind, too, that any encryption algorithm can be broken. The most basic component of

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

cryptography is mathematics. Every algorithm is based on a defined mathematical construct, so it is predictable to some degree. The major factors in cracking an algorithm are the amount of cryptographic knowledge required to design the attack and the amount of processing power needed to perform the attack.

In VPNs, AES has become widely supported. Many vendors offer modules that enable AES as an encryption method for VPNs and remote access; these modules usually do not require much configuration.

Secure Sockets Layer Many VPNs use **Secure Sockets Layer (SSL)** to encrypt tunneled information via the Web. SSL was developed by Netscape Communications Corporation as a way of enabling Web servers and browsers to exchange encrypted information. SSL gets its name from the following points:

- It uses public and private key encryption to create secure communications.
- It uses the sockets method of communication between servers and clients. (A socket is simply a combination of an IP address and a port number, such as 210.45.73.28:443.)
- It operates at the Network layer (Layer 3) of the OSI model. However, it can still provide a level of security that works between TCP and HTTP.

SSL is widely used on the Web; in fact, VPNs that use SSL can support only data that is exchanged by Web-enabled applications. For this reason, SSL is unlikely to replace IPsec completely as a security standard for VPNs, but as mentioned earlier, the popularity of Web-based business applications is making SSL VPNs very common.

An SSL session makes use of both symmetric and asymmetric keys. The asymmetric keys are used to start an SSL session, but symmetric keys are then generated dynamically for most of the transfer. It works like this:

1. The client connects to the Web server by using SSL.
2. The two machines authenticate each other and determine which ciphers and protocols to use. The client sends the server its preferences for an encryption method, the SSL version number, and a randomly generated number to be used later.
3. The server responds with the SSL version number, its cipher preferences, and its digital certificate. The digital certificate tells the client who issued it, a data range, and the server's public key. The server can ask the client for its own digital certificate at this point.
4. The client verifies that the date and other information on the digital certificate are valid. If so, the client generates a "pre-master" code and sends it to the server, using the server's public key. The client's digital certificate is also sent if the server requests one.
5. The server uses its private key to decode the pre-master code that the client sent. The server then generates a master secret key that both the client and server use to generate session keys—symmetric keys used during the SSL session to verify the integrity of data that is exchanged.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

6. The server and client exchange messages indicating that future messages will be encrypted with the session key. Next, the client and server send each other a separate, encrypted message that their side of the handshake is complete.
7. The handshake is completed, and the SSL session begins. Either side can renegotiate the connection at any time, and the process repeats if there is a lost connection, timeout, or user intervention, such as logging out.

SSL is a secure way to transmit data on the Web, such as credit card numbers used for online purchases. However, SSL can cause problems for firewalls that cannot interpret SSL data, in much the same way that firewalls can have trouble interpreting IPsec data that uses tunnel mode encryption.

Transport Layer Security **Transport Layer Security** (**TLS**) was designed to provide additional security for Internet communication. Although it is similar to SSL in operation and design, TLS adds some notable improvements. First, it uses a hashed message authentication code (HMAC) that provides extra security by combining the hashing algorithm with a shared secret key. Both parties must have the same shared secret key to authenticate the data. Second, TLS splits the input data in half, processes each half with a different hashing algorithm, and then recombines the data by using an exclusive OR (XOR) function. This method provides protection if one of the algorithms proves to be vulnerable.

TLS provides authentication and encryption mechanisms. It supports certificate-based authentication, as SSL does, and the handshake process is similar. For encryption, TLS can support a variety of ciphers, both symmetric and asymmetric, and several hashing algorithms, as summarized in Table 11-2.

Type of cipher	Cipher/algorithm
Public key (asymmetric)	Rivest-Shamir-Adleman (RSA), Diffie-Hellman, Digital Signature Algorithm (DSA)
Private key (symmetric)	Rivest Cipher (RC) 2, RC4, International Data Encryption Algorithm (IDEA), DES, 3DES, AES, and Camellia
Hashing algorithms	Message Digest (MD) 2, MD4, MD5, Secure Hash Algorithm (SHA), SHA-1, and SHA-256

© Cengage Learning 2014

Table 11-2 TLS cryptography

TLS/SSL uses symmetric cryptography for bulk encryption and asymmetric cryptography for authentication and key exchange. TLS runs below Application layer protocols, such as SMTP, FTP, HTTP, and LDAP, but above the UDP and TCP Transport layer protocols. This makes TLS useful for securing a variety of protocols. For the most part, TLS takes place automatically without user action, and makes it easier to administer firewalls and NAT. The downside is that TLS requires more processing power and administration time for managing certificates.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Most new VPN hardware devices support TLS/SSL, and it is available in current OSs. OpenVPN uses TLS/SSL to tunnel the entire network protocol stack.

Authentication

Authentication is the third core activity that VPNs perform to ensure the security of tunneled communication. Authentication is essential because network hosts that receive VPN communications need to know that the originator of the communication is an approved user of the VPN.

The type of authentication used in a VPN depends on the tunneling protocol. Many networks use IPsec to authenticate users; VPN participants establish an SA and exchange keys to authenticate one another. PPTP, which is used for dial-up access to a remote server, uses MS-CHAP, in which both computers exchange authentication packets. TLS/SSL uses certificate-based authentication. RADIUS is also commonly used for authentication services in VPN architecture.

VPNs use digital certificates to authenticate users and encryption to ensure that communications cannot be read even if they are intercepted in transit. Figure 11-9 illustrates the VPN's core activities of encapsulation, encryption, and authentication.

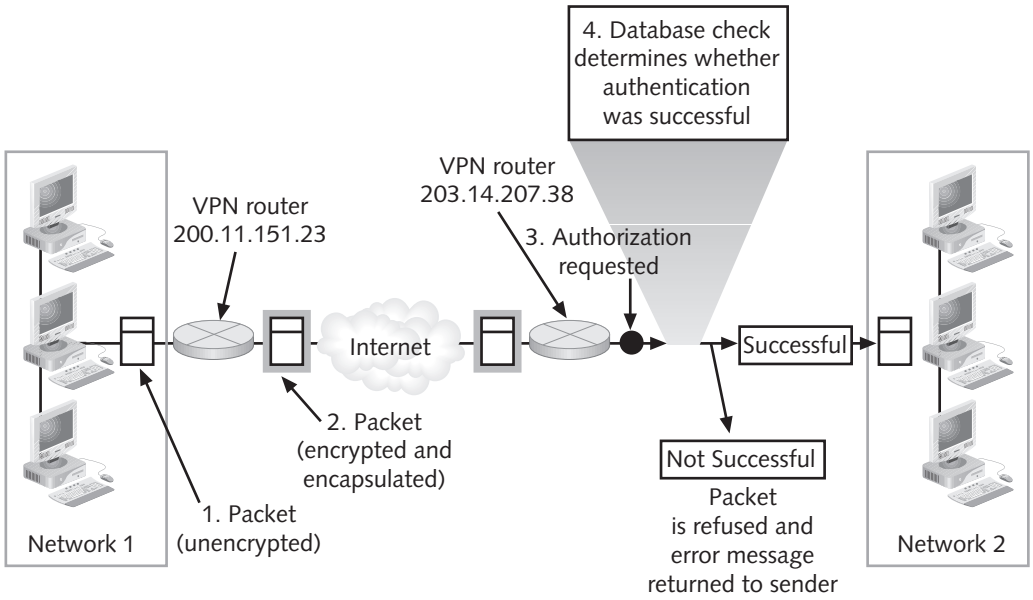

Figure 11-9 VPN core activities

In Figure 11-9, the following steps are basic to the authentication process:

1. The source computer transmits the unencrypted packet on internal network 1.
2. After the VPN router at 200.11.151.23 encrypts and encapsulates the packet, the packet passes through the gateway into the Internet.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

3. The VPN router at internal network 2 requests authentication.
4. A database check determines whether authentication is successful. If it is, the packet is allowed to reach its destination. Otherwise, an error message is returned.

Kerberos **Kerberos** is an authentication system developed at the Massachusetts Institute of Technology (MIT). The name refers to the three-headed dog in Greek mythology that guarded the gates of Hades. Kerberos authenticates the identity of network users by using a simple method called "authentication by assertion." The computer that connects to a server and requests services asserts that it is acting on behalf of an approved user of those services. Although this method sounds simple, the process by which computers communicate the assertion and response is not. Figure 11-10 shows how Kerberos works; the following list outlines what happens at each step:

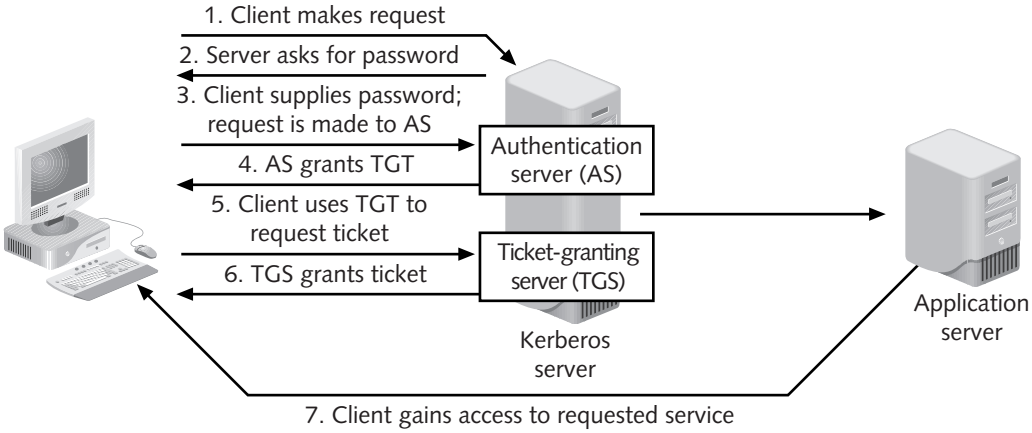

© Cengage Learning 2014

Figure 11-10 The Kerberos authentication process

1. A client requests a file or other service.
2. The server prompts the client for a username and password (user account credentials).
3. The client submits the credentials, and the request is sent to an authentication server (AS) that is part of the Kerberos system. The Kerberos AS is known as the **Key Distribution Center** (**KDC**). A Windows domain controller can serve as the authentication server.
4. The AS creates an encrypted code called a session key, based on the client's password and a random number associated with the service being requested. This session key is called a **ticket-granting ticket** (**TGT**).
5. The AS grants the TGT.
6. The client presents the TGT to the **ticket-granting server** (**TGS**). The TGS is also part of the Kerberos system, but it may not be the same server as the AS.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

7. The TGS grants the session ticket and forwards it to the server holding the requested file or service on behalf of the requesting client.
8. The client gains access to the requested service or file.

A major advantage of Kerberos is that passwords are not stored on the system, so they cannot be intercepted. A ticket is specific to a user and typically expires after a set period (usually eight hours). Kerberos does allow postdated tickets and renewable tickets, but users cannot modify tickets directly; they must request these flags when requesting their tickets.

Kerberos has a lower "network overhead" than a Public-Key Infrastructure (PKI) system, so you do not need to install a central server and manage a Kerberos system as you would a PKI system. This lower overhead is especially handy in a small network where only a few users need access to shared resources. Kerberos is also useful for single sign-on (SSO). Instead of having to sign on to each resource and keep track of all those credentials, users simply sign on to the Kerberos authentication services and have access to all resources they have permission to use.

Kerberos does come with a major concern, however. The AS or KDC is a single point of failure, so if it goes down and no backup servers can take over its role, no one can be authenticated. Administrators who use a single AS should take measures to ensure that authentication services can continue to function in the event of an AS failure. Other major security flaws have emerged in Kerberos, making it possible for attackers to run arbitrary code on KDCs, launch man-in-the-middle or DoS attacks, or cause the domain controller to shut down. Patches are available to address these vulnerabilities.

Examining VPN Design and Architecture

A VPN's **topology**—the way components in a network are connected—determines how gateways, networks, and clients are related to each other. As you will learn in the following sections, VPN topologies correspond to a basic network's physical and logical topologies. The three basic topologies are mesh, star, and hybrid VPNs.

Mesh Topology

In a **mesh configuration**, all participants in the VPN have Security Associations with one another. Two types of mesh arrangements are possible:

- *Full mesh*—Every subnetwork is connected to all other subnets in the VPN (see Figure 11-11). This topology is complex to manage and is best used with small VPNs.
- *Partial mesh*—Any subnet in the VPN may or may not be connected to the other subnets. This configuration offers more flexibility than a full-mesh arrangement.

The advantage of a mesh configuration is that each participant can establish VPN communications with all other participants. However, if a new LAN is added to the VPN, all other VPN devices have to be updated to include information about new users. The problem with mesh VPNs is the difficulty of expanding the network and updating every VPN device whenever a host is added.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

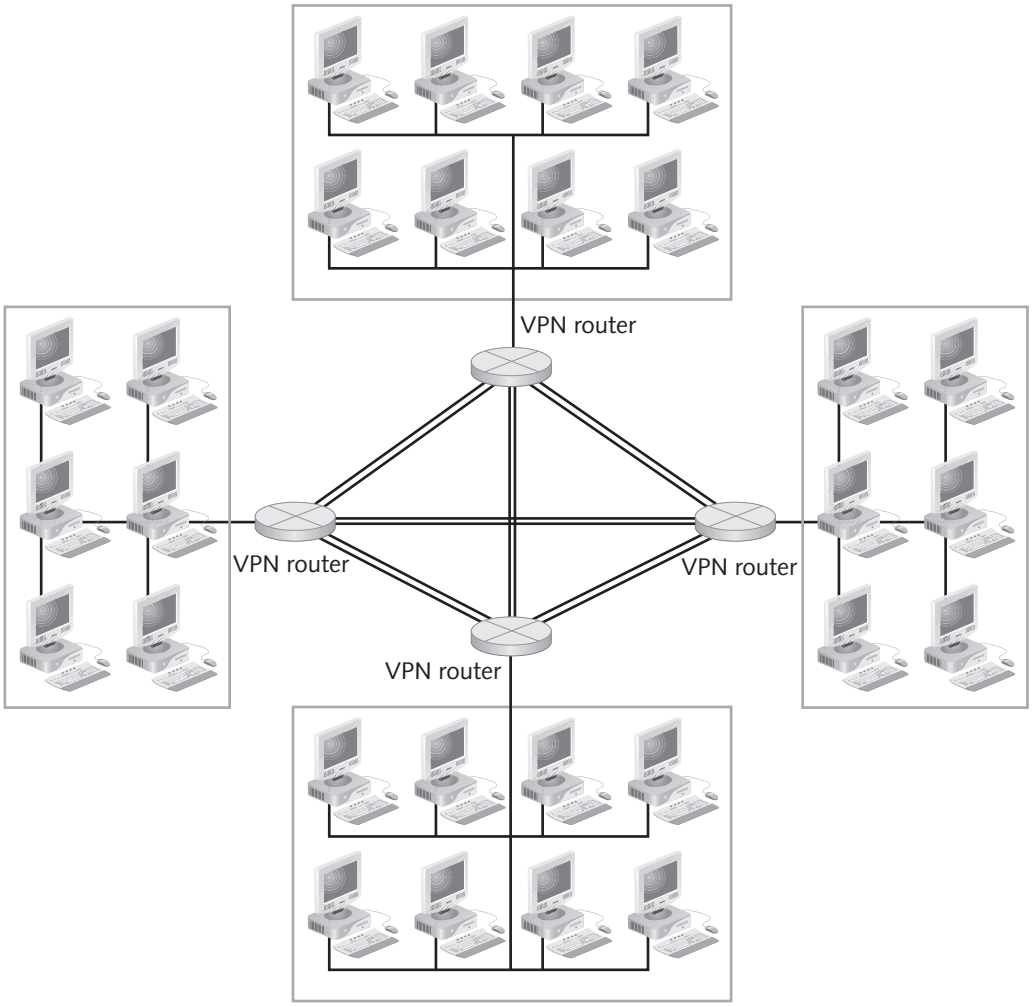

© Cengage Learning 2014

Figure 11-11 A full-mesh VPN configuration

Star Topology

In a **star configuration** (also known as a hub-and-spoke configuration), the VPN gateway is the hub, and other networks participating in the VPN are called rim subnetworks (see Figure 11-12).

In this configuration, separate SAs are made between the hub of each rim subnetwork. The central VPN router is at the organization's central office because most star VPNs have communications go through the office where the main IT staff is located. Any networks or computers that want to participate in the VPN need to connect only to the central server, not to any other systems in the VPN. This setup makes it easy to increase the VPN's size when more branch offices or computers are added. On the other hand, in star configurations, all communications flow in and out of a central router. This setup creates a single point of failure at the central router and can slow communication, especially if branch offices are far

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

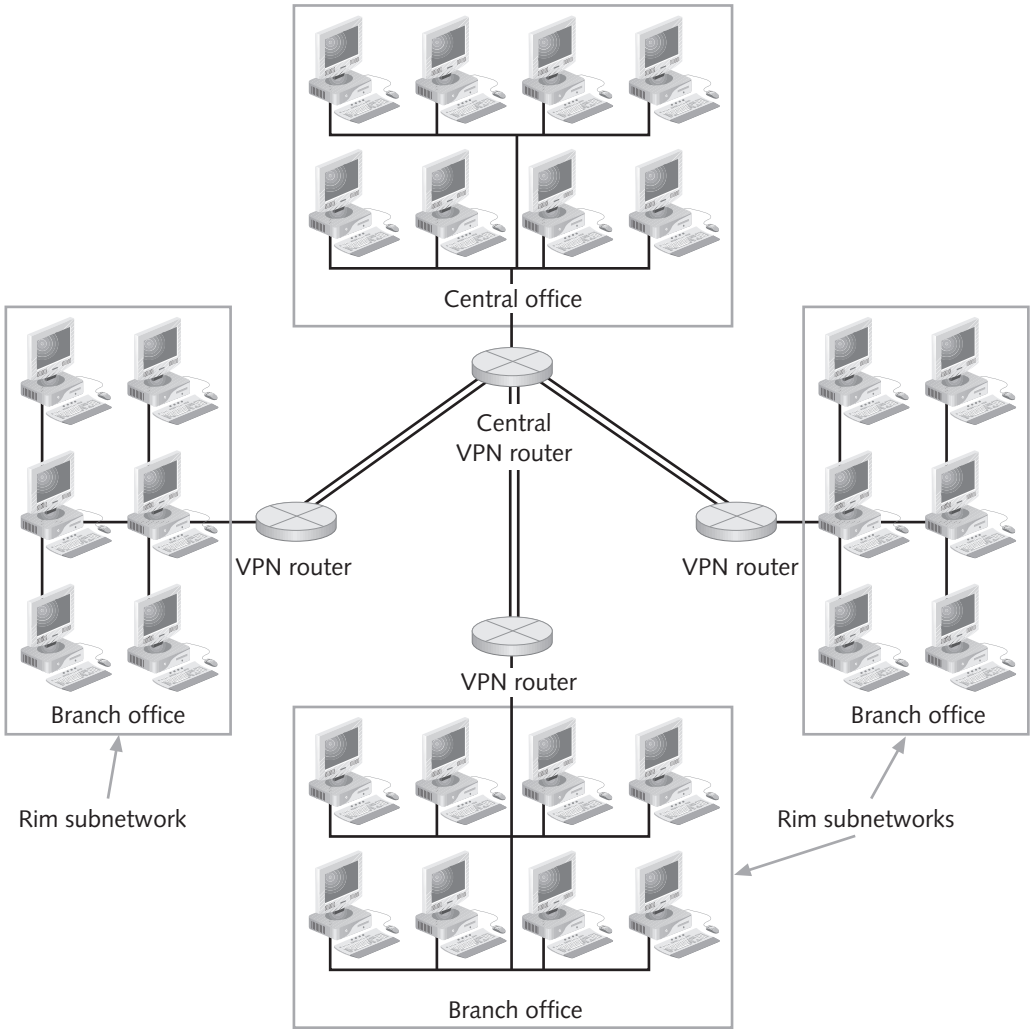

© Cengage Learning 2014

Figure 11-12 A star VPN configuration

apart. One solution is to use two or more routers at the central office. They can be configured for load balancing to improve performance and function as failover devices.

Hybrid Topology

As organizations with VPNs grow to include new computers and branch offices, they naturally evolve from a mesh or star configuration into a **hybrid configuration** that combines network topologies (see Figure 11-13). Because mesh configurations tend to operate more efficiently, the central core that links the network's most important branches should probably be a mesh configuration to provide fault tolerance. However, branch offices can be added as spokes that connect to a VPN router at the central office. A hybrid setup that combines these two configurations benefits from the star configuration's scalability and the mesh configuration's speed.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

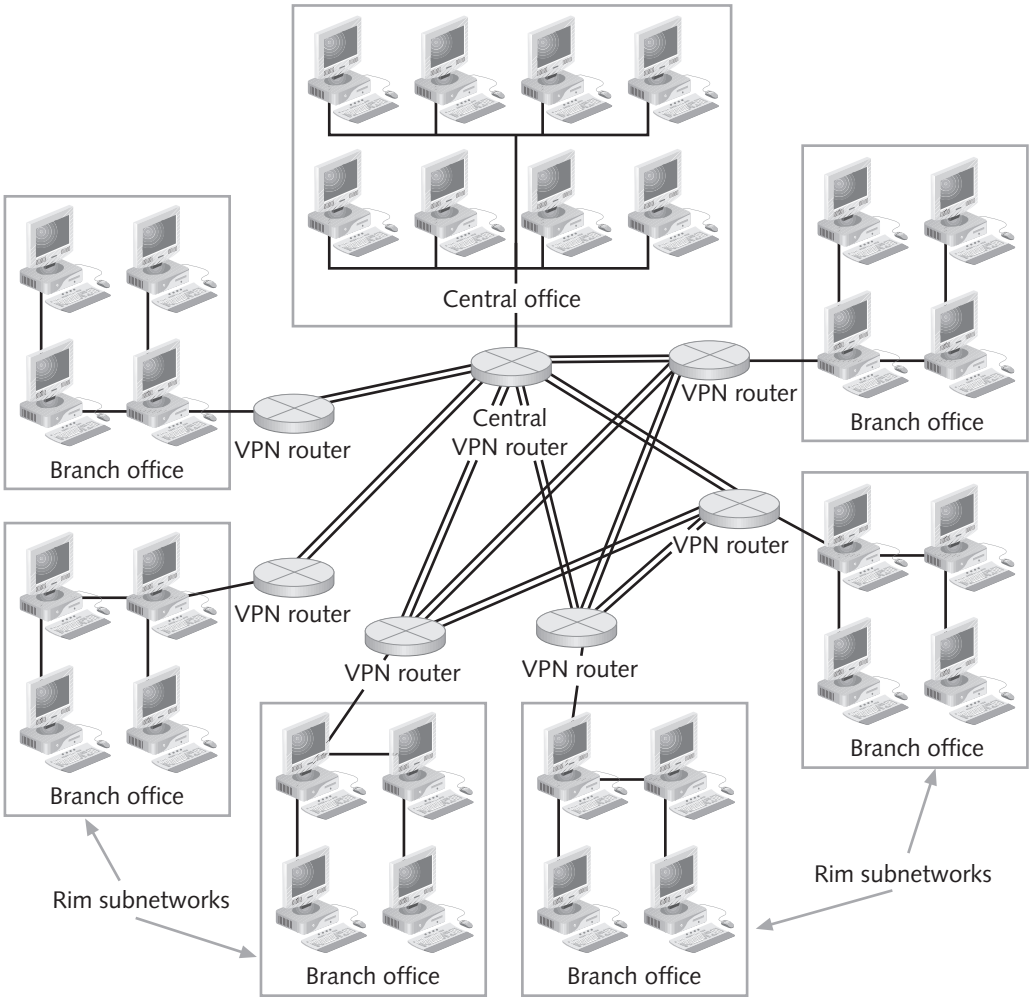

© Cengage Learning 2014

Figure 11-13 A hybrid VPN configuration

VPN Domains

To set up a VPN, you need to define a **VPN domain**: a set of one or more computers that VPN hardware and software handle as a single entity. The computers in a VPN domain use the VPN to communicate with another domain. With a firewall, a domain might be a set of networked computers grouped under a name, such as Office_Network.

Single and Multiple Entry Point Configurations The decision whether to have single or multiple entry points depends on whether your network has a site-to-site or client-to-site VPN configuration. Small networks that use VPNs typically have only site-to-site connections and often have **single entry point configurations**: All traffic to and from the network passes through a single gateway, such as a router or firewall (or both). In a single entry point configuration, the gateway must be a member of the VPN domain. In the

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

configuration shown in Figure 11-14, the VPN domain includes a group of computers in the internal network as well as the gateway.

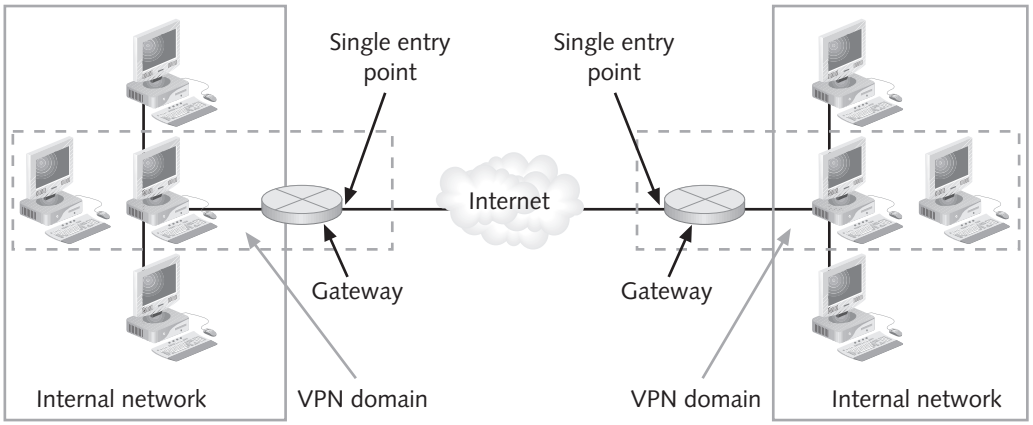

© Cengage Learning 2014

Figure 11-14 A single entry point configuration

In contrast, many large organizations have networks with several client-to-site connections. These connections require **multiple entry point configurations** in which multiple gateways are used, each with a VPN tunnel connecting a different location (see Figure 11-15). In a multiple entry point configuration, excluding the gateway from the VPN domain is important. If you do not exclude the gateway, all traffic to and from each gateway in the internal network is encrypted. This encryption reduces performance unnecessarily because you need to encrypt only the traffic from gateway to gateway.

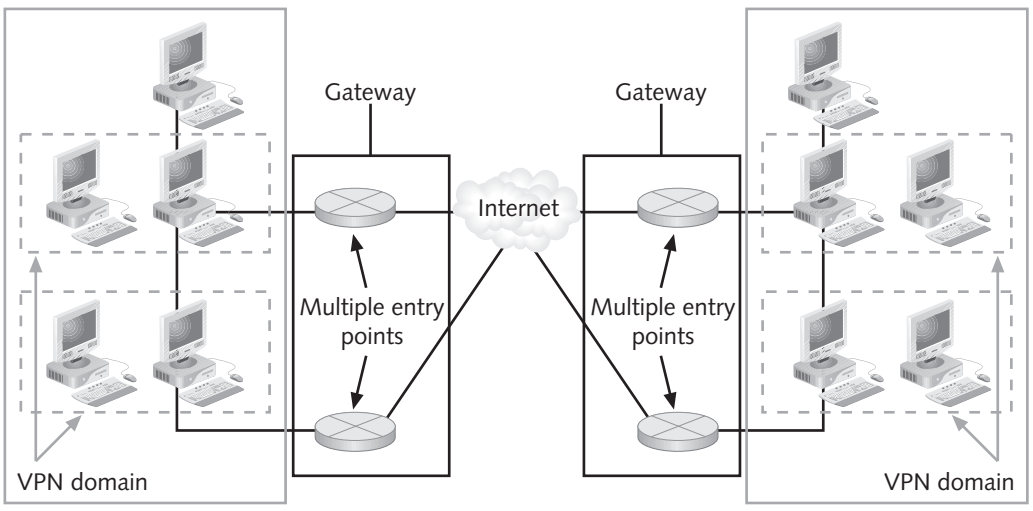

© Cengage Learning 2014

Figure 11-15 A multiple entry point configuration

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Preventing VPN domains from overlapping is also important; having multiple routes in routing tables could cause some traffic to be routed incorrectly or not at all because of duplicate IP addresses. If a router has multiple paths for directing packets, it might not respond correctly. This problem can be fixed easily by configuring routing tables correctly.

Using VPNs with Firewalls

Having a VPN does not reduce the need for a correctly configured firewall. You should always use a firewall as part of your VPN security design. Using a VPN with a firewall, however, requires careful planning and configuration. Several different configurations are possible, and each option has advantages and disadvantages, as you learn in this section.

One option is to install VPN software on the firewall. Several commercial firewalls include VPN components as an added option. As you can see in Figure 11-16, this configuration has a single point of entry into the network:

- The firewall allows outbound access to the Internet.
- The firewall prevents inbound access from the Internet.
- The VPN service encrypts traffic to remote clients or networks.

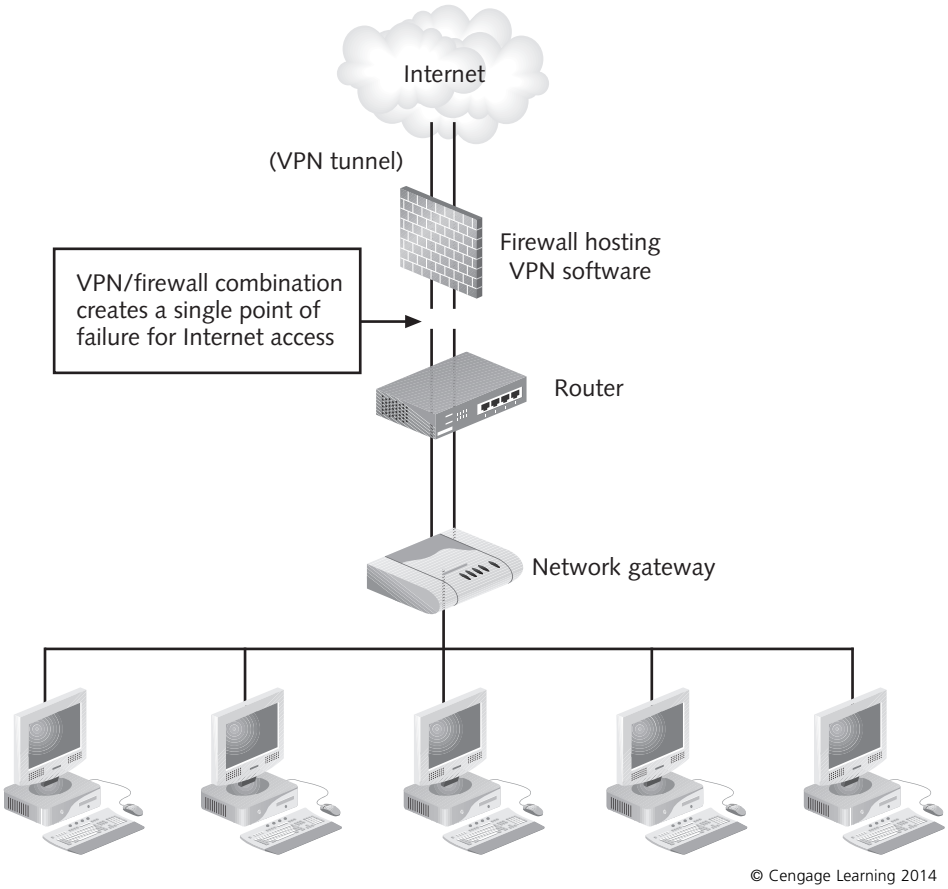

© Cengage Learning 2014

Figure 11-16 The VPN server on a firewall

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Putting the VPN on a firewall has the following advantages:

- You can control all network access security from one server.
- You have fewer computers to manage, meaning less chance of configuration mistakes.
- You can create and manage rules that apply to your VPN traffic with the same tools you already use to manage your firewall.

Installing the VPN on a firewall carries disadvantages as well:

- You have one server controlling all network access security. Any errors in configuring the VPN or firewall could leave your network open to attack.
- You must make sure to configure routes carefully so that traffic goes through the correct interfaces.
- Incorrect configuration of the firewall or VPN rules could allow traffic from the Internet to get past your security.
- Internet access and VPN traffic compete for resources on the server, so a more powerful computer might be necessary.

Another option, shown in Figure 11-17, is to set up the VPN parallel to your firewall inside the demilitarized zone (DMZ). Internal clients continue to point to the firewall as their default gateway and are unaware of the VPN connection. The firewall has a route to any networks accessible via the VPN server and instructs clients to send packets to the VPN server when appropriate.

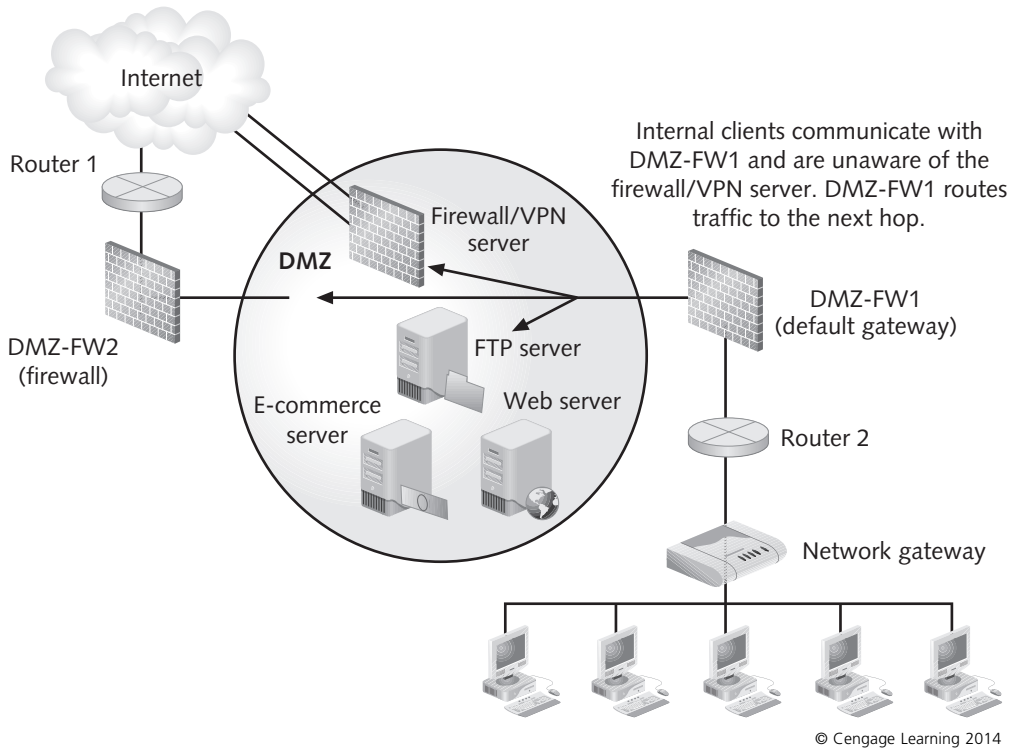

Figure 11-17 The VPN server parallel to a firewall

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

A DMZ can also be called a screened subnet or a perimeter network. If it is used for sharing files and access to company data with a business partner, a perimeter network is often referred to as an extranet.

Placing the VPN server parallel to the firewall has the following advantages:

- VPN traffic is not going through the firewall, so there is no need to modify firewall settings to support VPN traffic.
- This configuration can be scaled more easily. New VPN servers can be added without having to reconfigure the firewall.
- If the VPN server becomes too congested, you can add another server and distribute the load.

Placing a VPN server parallel to a firewall also includes the following disadvantages:

- The VPN server is connected directly to the Internet, making it an ideal target for attackers.
- If the VPN server becomes compromised, the attacker will have direct access to your internal network.
- The cost of supporting a VPN increases with the addition of new servers and extra support staff.

Another location for the VPN server is behind the firewall connected to the internal network. As shown in Figure 11-18, the VPN server is not accessible from the Internet. All packets must go through the firewall to reach the VPN server. As with the parallel configuration, you need to add a route to the firewall that redirects VPN traffic from internal clients to the VPN server. You also need to configure the firewall to pass encrypted VPN traffic directly to the VPN server.

Putting the VPN server behind the firewall has some advantages:

- The VPN server is completely protected from the Internet by the firewall.
- The firewall is the only device controlling access to and from the Internet.
- Network restrictions for VPN traffic are configured only on the VPN server, making it easier to create rule sets.

On the other hand, putting the VPN server behind the firewall has the following disadvantages:

- All VPN traffic must travel through the firewall, which increases congestion and latency.
- The firewall must handle VPN traffic from the Internet to the VPN server. Getting the firewall to pass encrypted VPN traffic to the VPN server could require advanced configuration.
- The firewall might not know what to do with IP protocols other than ICMP, TCP, and UDP. Supporting VPNs that use IP protocols, such as ESP packets for IPsec or GRE packets for PPTP, could be challenging.

If you terminate the VPN connection in front of the firewall, VPN traffic will be on the unprotected external network for a brief period before passing through the firewall.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

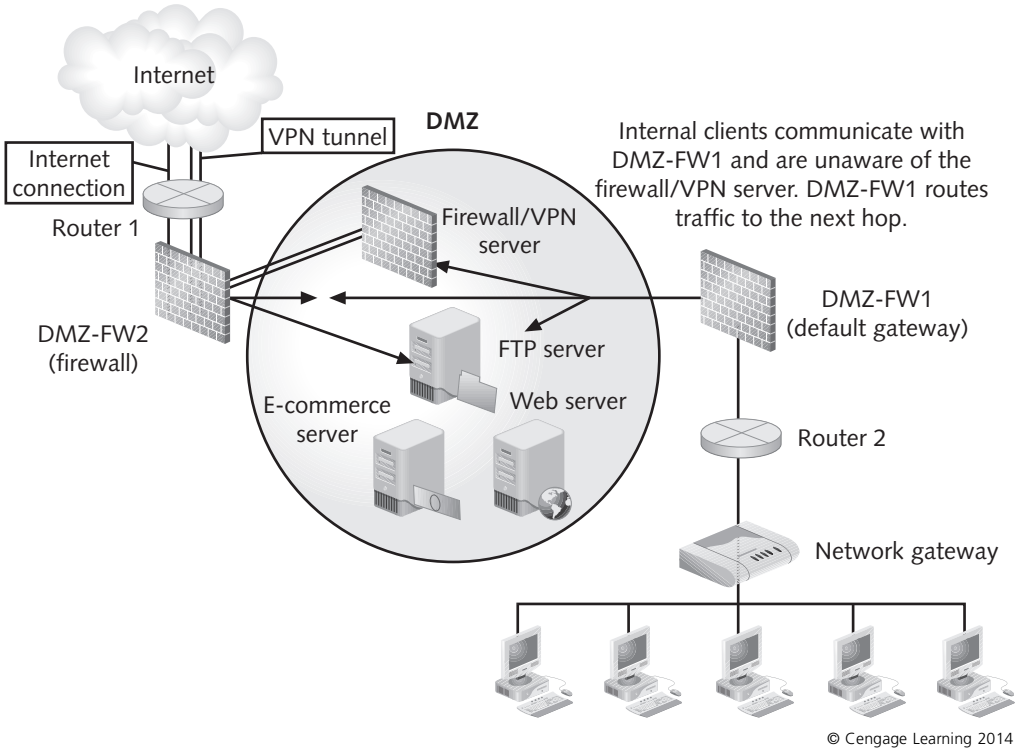

Figure 11-18 The VPN server behind a firewall

Adjusting Packet-Filtering Rules for VPNs

A perimeter firewall is configured with rules that filter packets the VPN sends or receives. These rules control intranet traffic to and from VPN clients and should be based on your organization's security policies. Packet filtering is based on the header fields of inbound and outbound packets. However, to configure your VPN and firewall combination, you need to set up packet-filtering rules that allow VPN traffic to pass through as needed. Packet filtering makes use of three IP packet header fields in particular:

- The source address is the 32-bit IP address of the host that sent the packet and generated the information in it.
- The destination address is the 32-bit IP address of the destination host expected to receive the packet and its information.
- The protocol identifier (protocol ID) is the number of the upper-level protocol to which the data in the packet belongs.

You can conduct packet filtering based on any or all of these header fields. Using the source or destination address, for example, you can block all packets from an address or range of addresses, and you can route allowed packets to a range of addresses. The protocol ID field is used to refer to protocols such as the following:

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- ICMP, protocol ID 1
- TCP, protocol ID 6
- UDP, protocol ID 17
- GRE, protocol ID 47
- ESP, protocol ID 50
- AH, protocol ID 51

The Internet Assigned Numbers Authority (IANA) maintains a list of all protocol IDs at *www.iana.org/assignments/protocol-numbers*.

PPTP Filters PPTP, the first widely supported VPN protocol, supports legacy authentication methods, does not require a PKI, and provides automated configuration. PPTP might be the best option when VPN connections must pass through a NAT server or firewall. NAT changes the internal computer's IP address to the address of the NAT device. For this reason, NAT is not compatible with many IPsec implementations. The problem comes from the IKE protocol that IPsec uses. IKE embeds the sending computer's IP address in its payload, and because the embedded address does not match the source address of the IKE packet (which is the NAT device's address), the receiving computer drops the packet. NAT traversal techniques have been developed to overcome this problem.

Currently, there is no single accepted NAT traversal method, but using UDP ports 500 and 4500 to negotiate UDP encapsulation-based addressing is a common method when IPsec/IKE is employed.

PPTP comes into play when older clients need to connect to a network through a VPN or when a tunnel must pass through a firewall that performs NAT but is not compatible with IPsec. For PPTP traffic to pass through the firewall, you need to set up packet-filtering rules that permit this communication. PPTP uses two protocols: TCP and GRE. A VPN server that has been configured to receive PPTP traffic listens for incoming connections arriving on TCP port 1723. It also needs to receive GRE packets (identified by protocol ID 47). Table 11-3 shows filtering rules you would use for your gateway if it had an IP address of 205.43.1.78 and a remote gateway with an IP address of 77.127.39.2.

Rule	Source IP	Destination IP	Protocol	Source port	Destination port	Action
1	Any	205.43.1.78	TCP	Any	1723	Allow
2	Any	205.43.1.78	Protocol ID 47	Any	Any	Allow

© Cengage Learning 2014

Table 11-3 PPTP packet-filtering rules

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

In Table 11-3, two rules are established. Rule 1 allows incoming PPTP connections from any computer to be received at the VPN server at 205.43.1.78 via port 1723. Rule 2 allows incoming traffic that uses protocol ID 47 (incoming traffic consisting of GRE packets).

For increased security, you could specify an option in the Action column such as "Drop all packets except those that meet this criteria," if your firewall enables you to make these distinctions. In your rule base, you must configure this option as any other rule, which means you specify parameters such as source, destination, protocol, and action to take (Deny, in this case).

L2TP and IPsec Filters If you use L2TP, you need to set up rules that permit IPsec traffic. You have to account for IKE using protocol ID 1701 and UDP on port 500. As listed previously, ESP uses protocol ID 50, and AH uses protocol ID 51. Table 11-4 shows the filter rules for a network gateway with an IP address of 205.43.1.78 and a remote gateway with an IP address of 77.127.39.2.

Rule	Source IP	Destination IP	Protocol	Transport protocol	Source port	Destination port	Action
1	Any	205.43.1.78	IKE	UDP	500	500	Allow
2	Any	205.43.1.78	ESP	Protocol ID 50	Any	Any	Allow
3	Any	205.43.1.78	AH	Protocol ID 51	Any	Any	Allow

© Cengage Learning 2014

Table 11-4 L2TP packet-filtering rules

In this example, Rule 1 covers inbound IKE traffic, which uses UDP port 500. Rule 2 allows traffic if you decide to use ESP, which requires a filter rule allowing protocol ID 50 packets. Rule 3 allows traffic if you use AH, which requires a rule for protocol ID 51. As with filters for PPTP, you achieve better security if you specify a cleanup rule in the Action column, such as "Drop all packets except those that meet this criteria."

Ensuring Client Security

Another critical aspect of deployment is enforcing security on the client side of the VPN tunnel. If a remote user's computer is compromised, the internal network is at risk. Client-side issues include whether to require clients to use a firewall and intrusion detection and prevention system (IDPS), and whether policies should be enforced on client computers before allowing remote users to authenticate to the internal network.

Remember that a VPN extends the corporate network by using public communication channels. When it is set up and configured correctly, a VPN can supply inexpensive, secure access; when set up incorrectly, a VPN is a tremendous security liability that can give intruders full access to your network. Failing to secure access to a remote network has legal implications, too. If your VPN connects to partner or vendor networks, and an intruder gains access to their network through your poorly configured VPN, your company could face litigation for

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

damages resulting from the security breach. This situation isn't likely, but it is possible. The point is that you must consider the security implications of any extension of your network.

There are several ways to increase VPN client security, from network configuration settings to third-party software solutions. For example, when a Windows client connects to a VPN server, administrators can configure the remote system to use the corporate network's default gateway for further Internet access. This option prevents **split tunneling** by the client, which results in multiple paths. One path goes to the VPN server and is secured, but an unauthorized and unsecured path permits users to connect to the Internet or another network while still connected to the corporate VPN (see Figure 11-19). Split tunneling leaves the VPN server and internal network vulnerable to attack.

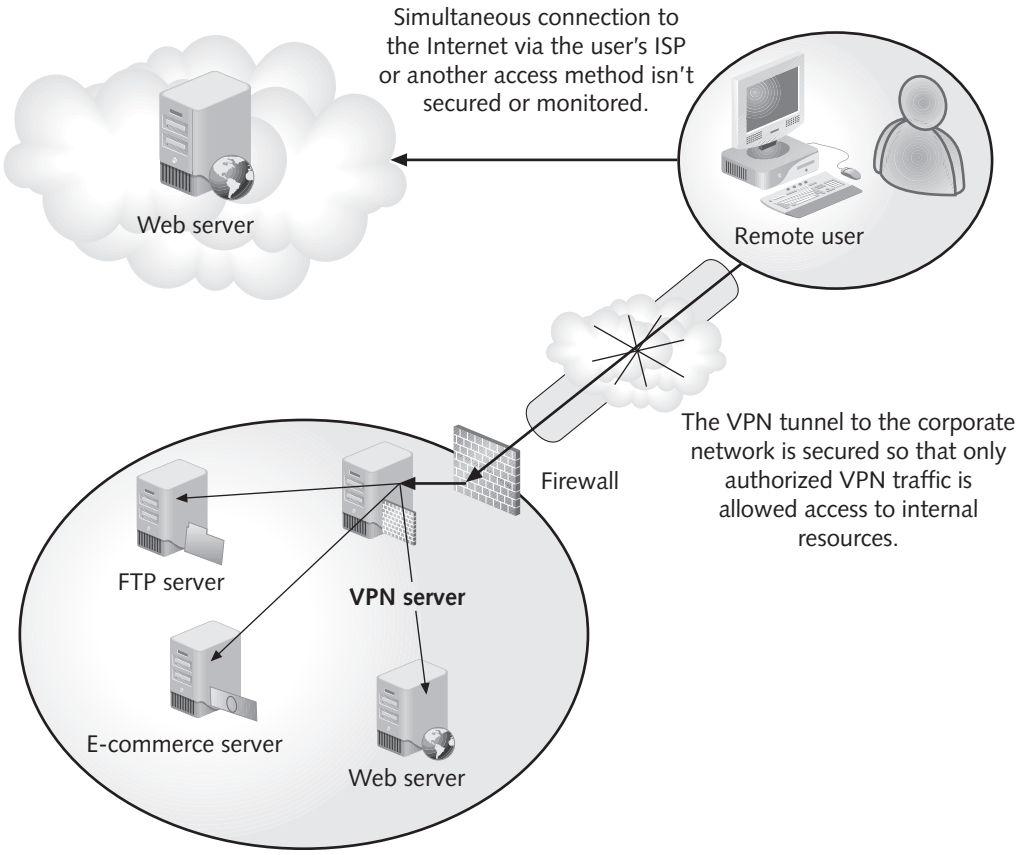

© Cengage Learning 2014

Figure 11-19 Split tunneling introduces a vulnerability

Small organizations that do not have the resources to enforce total control over clients should at least require users to sign an acceptable use agreement. Remote users should be required to maintain an Internet firewall and current virus protection to secure the computer against unauthorized use, and should be prevented from enabling split tunneling while connected to the network.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Auditing VPNs and Setting VPN Policies

VPN endpoints are vulnerable to many of the same viruses, Trojan programs, and spyware as internal network computers, and an infected remote computer can spread viruses or malicious code to computers on the internal network. Because remote clients are outside an administrator's realm of control, the responsibility for remote client security falls mostly on end users. You can take measures to mitigate these risks, however, such as including VPN policies in your overall security policy, using VPN quarantine procedures, logging VPN connections, and auditing remote access activity logs. These measures are discussed in the following sections.

Using VPN Quarantine

A quarantine network is a segment of the network with limited access to resources, especially internal resources. Quarantine networks are logical topology structures that are used for untrusted clients, such as business partners, vendors, and consultants, as well as remote access/VPN clients. Quarantined clients can be checked for policy compliance and then granted access to resources after compliance is verified, or a connection can be maintained in quarantine state for the duration of a session, depending on the client's access permissions and the nature of the connection (for example, a remote access/VPN client or a limited access consultant).

VPN quarantine was created to address the problem of remote clients not meeting an organization's security standards. Although a VPN provides encryption and encapsulation to secure access, it cannot check to make sure that clients have the latest updates or check for malicious software. A computer that connects to a network via a VPN can introduce malware; if the operating system or the antivirus software is not updated, it might contain software vulnerabilities that can be exploited. VPN quarantine provides a method to address these vulnerabilities by placing clients in virtual solitary confinement while they are checked for policy compliance.

Clients are subjected to preconnection and postconnection checks with custom scripts or a third-party add-on utility. The checks can examine the service pack version, software update status, and whether an approved antivirus program and personal firewall are running and updated on the client. Other requirements can be specified in the scripts. While these checks are running to verify compliance with the remote access policy, computers that attempt to connect are placed in a quarantine network. After computers are examined and pass the checks, the quarantine is lifted, and clients are allowed access to network resources. VPN quarantine is not a cure-all for security vulnerabilities, but it can help prevent computers with unsafe configurations from connecting to the internal network.

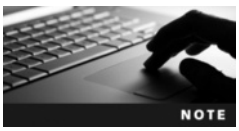

Quarantine is not used solely for remote access. It can be used to mitigate other threats, such as e-mail, which is a primary source of malware infections. Many corporate e-mail systems use an approach similar to quarantine for screening suspicious e-mails or e-mails with attachments. If an e-mail contains viruses or does not comply with policies, it is dropped, and the recipient might get an e-mail that explains the problem.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

To use VPN quarantine, you must have quarantine-compatible remote access clients and servers, resources for a quarantine network, an accounts database, and a quarantine remote access policy. VPN quarantine can use Windows authentication or a RADIUS server, but RADIUS is the preferred method.

Windows Server 2008 R2 and Threat Management Gateway 2010 support VPN quarantine scripts, and most other modern OSs can be configured to support quarantine procedures. Because VPN quarantine uses scripts and integrates with RADIUS servers, you can implement it on varied platforms easily. Refer to vendor instructions for procedures.

Logging VPN Activity

In a Microsoft network, the Routing and Remote Access Services (RRAS) server can log local machine events and record them in the server's System, Security, and Application logs. You can see these events in Event Viewer. Typically, the logs are used to notify administrators of unusual events or for troubleshooting. Event logs provide limited information about remote access, but specific logs give you more useful data about remote clients.

The RRAS server that hosts VPN services has additional logging capabilities beyond standard Event Viewer logs. It can also perform local authentication and accounting logging, which track remote access authentication attempts and use. The name of the applicable remote access policy is included with each connection attempt, making these logs useful for troubleshooting problems with remote access policies. You might need to enable Windows Authentication and Windows Accounting before you can configure what to log or where to store the logs.

Windows also provides support for RADIUS-based authentication and accounting logging. The RADIUS server stores logs in a separate file, which can be used for tracking remote access and use. You must enable RADIUS accounting and authentication before you can configure logging.

You must make sure that log files cannot be tampered with, maintain a reasonable retention policy, and ensure that log data is reviewed regularly. Storing a backup copy of log files on a separate computer is a good idea for sensitive data, and including log files in backups is a sensible precaution. Using a database with third-party log-file analysis tools can help you keep up with log file reviews. Using an integrated database that stores your remote access, firewall, and intrusion detection logs helps you compare logged events from all devices to get a clear picture of an event.

Auditing Compliance with VPN Policies

An organization's VPN or remote access policy defines standards for connecting to the network from any host or remote system. These policies must be integrated with an organization's overall security policies. Policies should be defined for different levels of restriction, such as what time of day access is allowed. You might also want to define tighter controls for business partners than you do for company employees working from home. Controls for administrators might be less restrictive; however, administrators should be required to have more secure passwords and to change them more often.

You should audit these policies to confirm that they are being enforced and that all users are complying with them. Enforcing these policies can be difficult with remote users, however.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Having strict policies for client-side OSs, configuration, and VPN client software makes this task easier, but configuring remote computers often requires long phone calls with the help of tech support, which can be frustrating and time consuming. One solution is to standardize the VPN client for remote users. That way, tech support staff need to learn how to support only one application, so there is less chance of errors.

In addition, test each client that will connect to the network to ensure smooth operation and help prevent security threats. After bugs are worked out and remote users can access the network, verify that everything is working according to the organization's policies and procedures. Remote user connections should be monitored for performance and capability to connect. Is the connection maintained during file transfers? How long does a normal file transfer take? Is an idle connection terminated after a specified time? You should work with a knowledgeable remote user to help determine a baseline for future auditing, testing, and troubleshooting.

Guidelines for VPN Policies

Incorporate the following best practices into your remote access and VPN policy:

- Plan the most secure deployment possible, and keep careful records of all changes.
- Use strong authentication methods.
- Require adequate password strength, length, and complexity, and require passwords to be changed frequently. Make sure that the password history setting is long enough to prevent the reuse of passwords.
- Have the remote access server use DHCP to assign addresses to remote clients or configure a static range of IP addresses on the VPN server so that it can assign addresses dynamically to remote clients.
- Log remote connections and use a centralized database or server for storing log files, if possible.
- Use VPN quarantine procedures, if available, to ensure policy compliance.
- Disable public peer-to-peer (P2P) file-sharing programs. If these programs are necessary for productivity, provide approved programs and ban all others.
- Preventing users from downloading and installing software is usually the safest route.
- Make sure that user accounts do not have full administrative access to their systems.
- Ensure strong encryption for data, especially passwords.
- Disable split tunneling.
- Require remote clients to be configured automatically to limit or eliminate user intervention. In general, you should prevent users from making security decisions.
- Disable or remove vulnerable protocols, such as Telnet, FTP, and rlogin.
- Use extra caution when configuring connectivity methods, especially wireless. Configure automatic disabling of wireless connectivity when users are connected directly or using another method of access.
- Prevent the use of removable storage devices, such as thumb drives or external hard drives, anywhere on the network.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Install personal firewall and antivirus programs on remote devices and configure them for automatic updates.
- Make sure that remote clients are self-defending, which means remote users cannot disable or bypass security measures. Whenever possible, measures should take place automatically without user intervention.
- Manage user audit information so that remote users receive policy updates and transmit their audit data before connecting.
- Conduct regular user training on security topics.

Try to strike the best balance between productivity and security in remote access and VPN policies. It is easy to err on the side of caution and lock a network down too tightly, which erases all the benefits of mobility. It is equally easy to err on the side of productivity and leave the network open to attack. Try to provide the best possible security within the guidelines of your security policy, yet give remote users enough freedom to be productive. Remember that overly stringent or complex policies are likely to be bypassed or ignored.

Chapter Summary

- A virtual private network uses encryption, encapsulation, and authentication to permit secure communications to travel over an unsecure transmission medium. VPNs have become popular because they allow organizations to use inexpensive Internet connections to conduct confidential business instead of having to lease dedicated communication lines.
- VPNs are created by hardware or software components at each end of the connection. Hardware devices are designed for performance but are more expensive than software VPNs. Software VPNs allow more flexibility of functional options.
- Analysis of business needs is an important first step in deploying VPNs. Excessive or inadequate security can result in unnecessary expense or risk. The business processes that require transmission security could involve remote employees, business partners, or branch office connectivity. A clear understanding of the business requirements is needed to design an effective and efficient VPN.
- A VPN's three core activities are encapsulation, encryption, and authentication. Encapsulation is the process of placing the original packet into another packet. Thus, the original source and destination addresses are protected. Tunneling protocols such as PPTP, L2TP, and SSTP are used to perform encapsulation.
- Tunneling protocols are associated with specific cryptographic methods. When PPTP is used, MPPE is the most commonly used encryption scheme. IPsec encryption is associated with L2TP.
- The authentication methods you use are dependent on the VPN type. For example, CHAP is used with PPTP and Authentication Header is used with L2TP.
- VPNs can use a full-mesh, partial-mesh, star, or hybrid topology. Depending on the size of the organization and the number of locations used to initiate connections, single or multiple entry points may be required.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- VPN protocols use specific ports. Packet-filtering rule sets must be configured to allow VPN traffic, except when a VPN endpoint is placed parallel to a firewall.
- VPNs must be managed in accordance with security policies. These security policies can be enforced on remote access systems through quarantine methods, which ensure that the remote system is secure before allowing it to connect to the internal network. Log auditing and policy enforcement must be consistent.

Key Terms

Authentication Header (AH) An IPsec protocol that provides authentication of TCP/IP packets to ensure data integrity.

client-to-site VPN A type of VPN connection that makes a network accessible to remote users who require dial-up access; also called a *remote access VPN*.

Encapsulating Security Payload (ESP) An IPsec protocol that encrypts the header and data components of TCP/IP packets.

encapsulation The process of enclosing a packet within another packet that has different IP source and destination information to ensure a high degree of protection.

Generic Routing Encapsulation (GRE) A nonproprietary tunneling protocol that can encapsulate a variety of Network layer protocols.

hybrid configuration A VPN configuration that combines characteristics of the mesh and star configurations.

Internet Key Exchange (IKE) A form of key exchange used to encrypt and decrypt data as it passes through a VPN tunnel. IKE uses tunnel method encryption to encrypt and then encapsulate packets for extra security. See also *tunnel method encryption*.

Internet Protocol Security (IPsec) A set of standard procedures that the Internet Engineering Task Force (IETF) developed for enabling secure communication on the Internet.

Internet Security Association Key Management Protocol (ISAKMP) An IPsec-related protocol that enables two computers to agree on security settings and establish a Security Association so that they can use Internet Key Exchange. See also *Internet Key Exchange (IKE)* and *Security Association (SA)*.

IPsec driver Software that handles the tasks of encrypting, authenticating, decrypting, and checking packets in an IPsec connection.

Kerberos An IETF standard for secure authentication of requests for resource access.

Key Distribution Center (KDC) A Kerberos component that holds secret keys for users, applications, services, or resources; KDC creates and distributes session keys by using symmetric cryptography.

Layer 2 Tunneling Protocol (L2TP) A tunneling protocol derived from two older protocols (Cisco's L2F and Microsoft PPTP L2TP) that encapsulates PPP packets and is usually combined with IPsec for improved security.

mesh configuration A VPN configuration in which all participants in the VPN are connected to one another. This configuration is commonly arranged as a full-mesh or partial-mesh setup.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

multiple entry point configuration A type of VPN configuration in which multiple gateways are used, each with a VPN tunnel connecting a different location.

Point-to-Point Tunneling Protocol (PPTP) A tunneling protocol used for dial-up access to a remote server.

Secure Sockets Layer (SSL) A protocol developed by Netscape Communications Corporation as a way of enabling Web servers and browsers to exchange encrypted information.

Security Association (SA) A designation for users, computers, or gateways that can participate in a VPN and encrypt and decrypt data by using keys.

side channel attack An attack method that exploits vulnerabilities of underlying hardware systems that leak data instead of exploiting vulnerabilities in a cryptographic algorithm.

single entry point configuration A VPN configuration in which all traffic to and from the network passes through a single gateway, such as a router or firewall.

site-to-site VPN A VPN that uses hardware devices, such as routers, to connect two networks; also called a *gateway-to-gateway VPN*.

split tunneling The term used to describe multiple paths. One path goes to the VPN server and is secured, but an unauthorized and unsecured path permits the user to connect to the Internet or some other network while still connected to the corporate VPN.

star configuration A VPN configuration in which a single gateway is the "hub" and other participating networks are considered "rim" networks.

ticket-granting server (TGS) The part of the KDC that creates and distributes session keys used by clients to access resources. See also *Key Distribution Center (KDC)*.

ticket-granting ticket (TGT) A digital token sent from the Authentication Server to the client. The client presents the TGT to the TGS to obtain a session key to access the resource. See also *Key Distribution Center (KDC)* and *ticket-granting server (TGS)*.

topology The way in which systems in a network are connected to one another.

Transport Layer Security (TLS) A protocol designed to secure Internet traffic. Although it has not replaced SSL yet, it offers improvements and is used more widely than SSL for a variety of applications.

tunnel The connection between two endpoints in a VPN.

tunnel method encryption A method of key exchange that encrypts both the header and data components of a packet and encapsulates it within a new packet that has a different header.

Unified Threat Management (UTM) A term for products that integrate a variety of security features into a single application, device, or product. These features include VPN and remote access services, firewalls, intrusion detection and prevention functions, and management consoles.

virtual private network (VPN) A cost-effective way for networks to create a secure private connection using public lines (usually the Internet). VPN endpoints establish connections (tunnels) to transmit and receive data, and then tear down the connections when they are no longer needed. Combinations of encryption, authentication, and encapsulation help ensure the confidentiality, privacy, and integrity of information.

VPN appliance A hardware device designed to terminate VPNs and join networks.

VPN client A router or an OS that initiates a connection to a VPN server.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

VPN domain A group of one or more computers that the VPN hardware and software handle as a single entity. This group uses the VPN to communicate with another domain.

VPN protocols Standardized communication settings that software and hardware use to encrypt data sent through a VPN.

VPN quarantine A method to address the problem of remote clients not meeting an organization's security standards. Quarantine places remote clients in a secured area while they are checked to ensure that software updates and current patches have been applied, antivirus software has been installed and updated, and other policies are in compliance.

VPN server A computer configured to accept VPN connections from clients.

Review Questions

1. What client-side issues do you need to consider when planning a VPN deployment? (Choose all that apply.)
   a. whether to require the client to use a firewall
   b. the organization's current growth rate
   c. how policies should be enforced on the client computer
   d. the cost of equipment that employees need to buy

2. In a mesh topology, all participants in the VPN have _____ with one another.
   a. tunnels
   b. SAs
   c. static routes
   d. trusts

3. What is a main disadvantage of mesh VPNs?
   a. They are not reliable.
   b. There is a lack of confidentiality among peers.
   c. They are difficult to enlarge or change.
   d. The equipment must be the same at all sites.

4. Putting a VPN on the firewall has which of the following disadvantages? (Choose all that apply.)
   a. There are more computers to manage.
   b. Only one server controls security, so any configuration errors leave the network open to attack.
   c. Internet access and VPN traffic compete for resources on the server.
   d. VPN traffic is not encrypted.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

5. A VPN server configured to receive PPTP traffic listens for incoming connections on port _____ and needs to receive GRE traffic identified by protocol ID _____.
   a. UDP 1443, 17
   b. TCP 1723, 47
   c. UDP 3349, 443
   d. UDP 1723, 47

6. Which protocols and ports must be allowed to pass when you are using L2TP and IPsec? (Choose all that apply.)
   a. protocol ID 50
   b. UDP 500
   c. TCP 50
   d. protocol ID 1701

7. AH uses protocol ID _____.
   a. 50
   b. 171
   c. 500
   d. 51

8. The VPN connection through which data passes from one endpoint to another is called a(n) _____.
   a. gateway
   b. extranet
   c. tunnel
   d. transport

9. A group of authentication and encryption settings that two computers negotiate to set up a secure VPN connection is called which of the following?
   a. protocol
   b. Security Association (SA)
   c. handshake
   d. key exchange

10. What makes a VPN a cost-effective option?
    a. Computers can use the same hardware and software.
    b. It requires no administrative configuration to set up or maintain.
    c. Many VPN applications are available as shareware or freeware.
    d. VPNs use public Internet and ISP connections.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

11. IPsec provides for what security activity to take place before data is encrypted or transmitted?
    a. encapsulation
    b. authentication
    c. establishment of a Security Association (SA)
    d. application of security policy settings
12. Which of the following is an advantage of using a star VPN configuration?
    a. It is easier to increase the VPN's size.
    b. Fewer VPN hardware or software devices are required.
    c. Only the VPN server at the center or "hub" needs to be updated.
    d. All participants can communicate with all other participants.
13. Because of an increase in the use of Web-based business applications, there has been an increase in _____-based VPNs.
    a. SSL
    b. IPsec
    c. L2TP
    d. PPTP
14. Which of the following is *not* a best practice for VPN client management?
    a. Enable split tunneling.
    b. Disable FTP.
    c. Disable Telnet.
    d. Enable VPN quarantine.
15. Which of the following IP addresses is most likely to be the source IP address of an encapsulated VPN packet?
    a. 150.80.26.59
    b. 172.30.78.45
    c. 11.17.5.210
    d. 210.240.255.48

Hands-On Projects

Hands-On Project 11-1: Creating a Domain Group for VPN Remote Access

Time Required: 10 minutes

Objective: Create a user and a group for VPN remote access.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Description: To leverage the security of Active Directory in VPN management, you create a group that is given VPN access rights and you create a test user account that becomes a member of the group.

To complete the hands-on projects in this chapter, you must first complete the Chapter 10 hands-on projects.

The projects in this chapter are designed for three computers: the previously installed Windows Server 2008 domain controller, TMG running on another Windows Server 2008 computer, and the previously installed Windows 7 computer. IP addresses given in these projects are designed for one group of three computers. The instructor should assign each group a set of first octets that are different from those shown in the projects. For example, in the projects, the IP address ranges are 192.*x.y.z* and 10.*x.y.z*. The next group of three computers could be assigned IP address ranges of 193.*x.y.z* and 11.*x.y.z*, and so forth.

1. Start the Windows Server 2008 domain controller. You will create a group for remote access users.
2. Start Active Directory Users and Computers. In the Users folder, create a global security group named **Remote Client Access Users**.
3. In the Users folder, create a user account with the following configuration:
   - First name: **Remote**
   - Last name: **User**
   - User logon name: **ruser**
   - Password: **Pa$$word**
   - Uncheck the **User must change password at next logon** box.
   - Check the **Password never expires** box.
4. Make ruser a member of the Remote Client Access Users group.
5. Leave the server running for the next project.

Hands-On Project 11-2: Configuring Threat Management Gateway as a VPN Server

Time Required: 15 minutes

Objective: Configure a VPN server.

Description: In the previous project, you created the domain group to be used to allow remote access. In this project, you configure the VPN server to permit that group VPN access and define IP addresses on the internal network that will be issued to remote clients.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

1. Start the Threat Management Gateway (TMG) server that you created in the Chapter 10 hands-on projects.
2. Click **Start**, click **All Programs**, click **Microsoft Forefront TMG**, and click **Forefront TMG Management**.
3. Click **Remote Access Policy (VPN)**, and then click the **VPN Clients** tab if necessary. On the Tasks tab, click **Configure VPN Client Access**. Click **Add**, click **Advanced**, and click the **Locations** button. If necessary, enter the Team*x* administrator username and password in the Windows Security box. (In Team*x*, substitute your team number for *x*.) Expand **Team*x*.net**, click the **Users** folder, and click **OK**. In the Name line, verify that **Starts with** is selected, type **Remote** in the box, and then click the **Find Now** button. You should see Remote Client Access Users in the Search results box. See Figure 11-20. Click **OK**, and then click **OK** in the Select Groups window. In the VPN Clients Properties window, click **OK**.

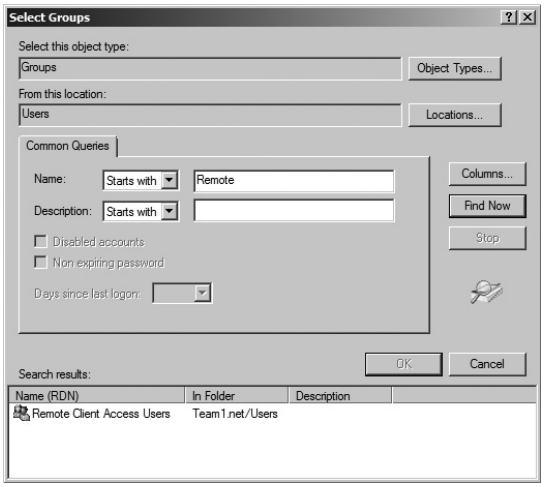

Used with permission of Microsoft Corporation

Figure 11-20 Finding the Remote Client Access Users group

4. To create a group of IP addresses to assign to VPN clients, click **Networking** in the left frame. Double-click **Internal** in the middle frame. Click the **Addresses** tab. Select the existing range, click **Edit**, and modify the Start address to **10.0.0.11**. The End address should remain as 10.255.255.255. Click **OK**. In the Internal Properties window, click **OK**.
5. Click **Remote Access Policy (VPN)** in the left frame. In the middle frame, click **Configure Address Assignment Method**. Verify that the **Static address pool** option button is selected, and click **Add**. Click the **Select the server** box, and then select your TMG server.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

6. In the Start address text box, type **10.0.0.1**. In the End address text box, type **10.0.0.10**. Click **OK**.
7. If necessary, click **Use the following network to obtain DHCP, DNS and WINS services**, and then select **Internal** from the menu.
8. Click the **Advanced** button. Click the **Use the following DNS server addresses** option button, and enter **10.0.0.125** as the IP address of your Windows Server 2008 domain controller in the Primary box. Leave the Backup box blank. Click **OK**.
9. Click the **Authentication** tab, and verify that **Microsoft encrypted authentication version 2** (**MS-CHAPv2**) is selected.
10. Click **OK**. In the top section of the Forefront TMG window, click the **Apply** button. Click **Apply** again. Click **OK**.
11. Leave your systems running for the next project.

Hands-On Project 11-3: Accessing the Internal Network from a Remote VPN Client

Time Required: 40 minutes

Objective: Use a VPN to connect to an internal network.

Description: In this project, you put the VPN to use. You create a VPN connection object on the remote client and connect to the TMG server. Then you create a firewall rule on TMG to allow the VPN clients access to the internal network. Once connected to the internal network, you map a drive on your Windows Server 2008 domain controller and place content on it.

1. Start Windows 7 and log in as an administrator.
2. Configure the Windows 7 IP settings. Enter an IP address of **192.168.1.120**, a subnet mask of **255.255.255.0**, a default gateway of **192.168.1.110**, and a DNS server address of **10.0.0.125**. You are now configured as an Internet client on the same network as the public side of the TMG server.
3. Click **Start**, and click **Control Panel**. If necessary, select **Small icons** from the View by menu in the upper-right corner, and click **Network and Sharing Center**.
4. Click **Set up a new connection or network**. Select **Connect to a workplace**, and click **Next**.
5. Click **Use my Internet connection (VPN)**. Click **I'll set up an Internet connection later**.
6. In the Internet address text box, type **10.0.0.111**, which is the IP address of your TMG server. In the Destination name text box, type **Connection to corporate headquarters**, and click **Next**.
7. In the User name text box, type **ruser**. In the Password text box, type **Pa$$word**. Click the **Remember this password** box. In the Domain text box, type **Team*x***. Click **Create**. Click **Close**.
8. In the Network and Sharing Center window, click **Connect to a network**. In the Set Network Location window, click **Work network**. Click **Connection to corporate**

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

headquarters, and click **Connect**. Click **Properties**, and click the **Security** tab. From the Type of VPN menu, select **Point to Point Tunneling Protocol (PPTP)**. Notice that Require encryption is selected. Examine the content of the other tabs, and then click **OK**. If necessary, type **ruser** in the User name box. In the Password box, type **Pa$$word**. If necessary, type **Team*x*** in the Domain box. Click **Connect**. This attempt should fail.

9. Log on to TMG as an administrator. If necessary, open the Forefront TMG Management console. In the left frame, click **Remote Access Policy (VPN)**. In the right frame, click **Configure VPN Client Access**. In the VPN Clients Properties window, click the **General** tab, and click the **Enable VPN client access** box. In the Maximum number of VPN clients allowed text box, type **5**, and click **Apply**. In the Microsoft Forefront Threat Management Gateway window, click **OK**. Examine the content of the other tabs. Be sure that PPTP is enabled on the Protocols tab. Click **OK**. In the Forefront TMG window, click **Apply**. Click **Apply**, and then click **OK**.
10. Return to Windows 7, and attempt to connect to the internal domain controller, as you did in Step 8. This attempt should succeed. Open a command prompt, type **ipconfig**, and press **Enter**. You should see your network adapter with its 192.168.1.120 address and a virtual PPP adapter with an IP address in the 10.0.0.0/8 network, which was issued to you from the range that you configured on TMG. This address allows you to communicate on the internal network.
11. Log on to the Windows Server 2008 domain controller. Click **Start**, click **Computer**, and double-click **Local Disk (C:)**.
12. Right-click any white space in the right frame, select **New**, and select **Folder**. Name the folder **Target**. Right-click **Target**, click **Properties**, click the **Sharing** tab, click **Advanced Sharing**, and click the **Share this folder** box.
13. Click **Permissions**, verify that **Everyone** is selected, and click the **Full Control** box in the Allow column. Click **OK**. Click **OK**.
14. Click the **Security** tab, and click **Edit**. In the Group or user names box, select **Users (TEAM*x*\Users)**. Click the **Full Control** box in the Allow column. Click **OK**, and click **Close**.
15. Return to Windows 7. At a command prompt, type **net use V:\\Team*x*SRV\Target** and press **Enter**. This attempt should fail. You are connected to the TMG server and you have a virtual network interface address on the internal network, but the firewall on TMG is keeping you out.
16. Return to TMG. In the Forefront TMG window, right-click **Firewall Policy** in the left frame. Click **New**, and click **Access Rule**. In the Access rule name text box, type **Allow VPN Access** and click **Next**. Click the **Allow** option button, and click **Next**. In the This rule applies to menu, select **All outbound traffic**, and click **Next**. Select **Do not enable malware inspection for this rule**, and click **Next**. Click **Add**, expand the **Networks** node, select **VPN Clients**, click **Add**, click **Close**, and click **Next**. Click **Add**, expand the **Networks** node, select **Internal**, click **Add**, click **Close**, and click **Next**. Accept the default user set of All Users, click **Next**, and click **Finish**.
17. In the Forefront TMG window, click **Apply**, click **Apply**, and click **OK**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

18. Repeat Step 15. Now that you have been allowed on the internal network, the command is successful. Change to the V: drive by typing **V:** at the command line and pressing **Enter**.
19. Create a file on the internal server from your remote Windows 7 system. Type **copy con January_Sales.txt** and press **Enter**. Type **January sales have been slow but are expected to pick up.** and press **Enter**. Type **Ctrl+Z** and press **Enter**.
20. Return to the Windows Server 2008 domain controller, and look in C:\Target. The new file should be present.

The command line was used to map a remote drive and create content on the remote server because it is an interesting process. However, a more common way to map a drive is to click Start, right-click Computer, and select Map network drive.

21. Leave your systems running for the next project.

Hands-On Project 11-4: Exploring TMG

Time Required: 60 minutes

Objective: Understand the features of TMG.

Description: In this project, you explore areas of TMG that you have not used in the previous projects. TMG provides much more than just proxy and VPN services. Sophisticated features provide intrusion detection and prevention, firewall functions, malware detection and prevention, and detailed logging functions. You explore these features in this project.

1. To begin, your systems should be in the same configuration as they were at the end of Hands-on Project 11-3. You should have a PPTP VPN connection between the Windows 7 computer and the Windows Server 2008 domain controller through the mapped drive V:, which points to \\Team*x*SRV\Target. Log on to TMG.
2. If necessary, open the Forefront TMG console. Click **Dashboard** in the left frame. Your result should look similar to Figure 11-21. Arrow 1 shows that one VPN client is connected, arrow 2 shows that there have been no violations of the Web access policy, arrow 3 shows that malware and intrusion detection (Network Inspection System) have not been updated with definitions, and arrow 4 shows the services running on TMG.
3. In the left frame, click **Monitoring**. If necessary, click the **Sessions** tab, where you should see the listing of the VPN client authenticated by ruser.
4. Click the **Alerts** tab to see information, warnings, and alerts. In Figure 11-22, you can see a warning about network activity that could be the result of IP spoofing.
5. In the left frame, click **E-Mail Policy**. Examine the content of the three tabs: E-Mail Policy, Spam Filtering, and Virus and Content Filtering.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

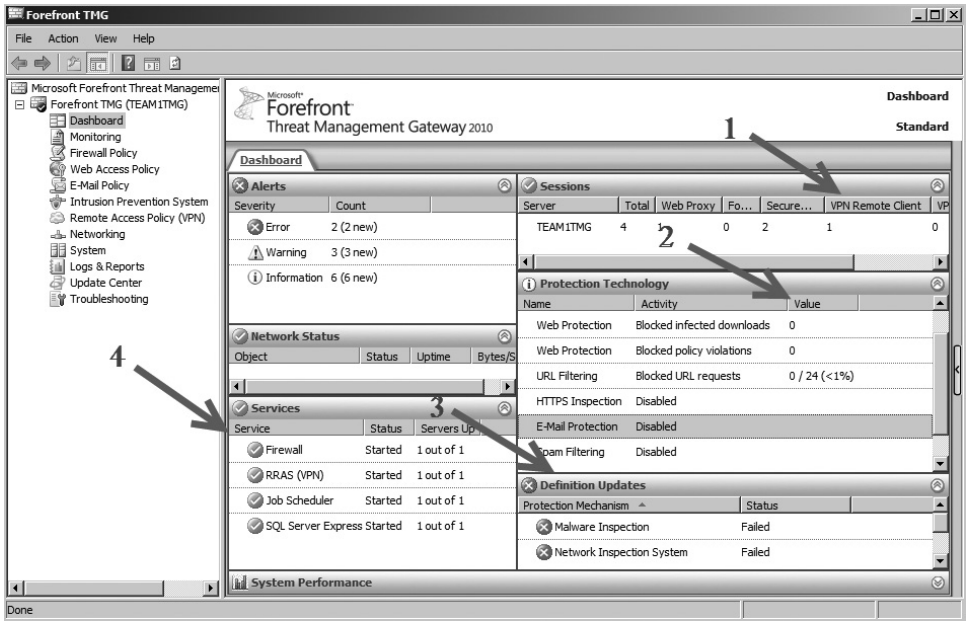

Used with permission of Microsoft Corporation

Figure 11-21 TMG Dashboard

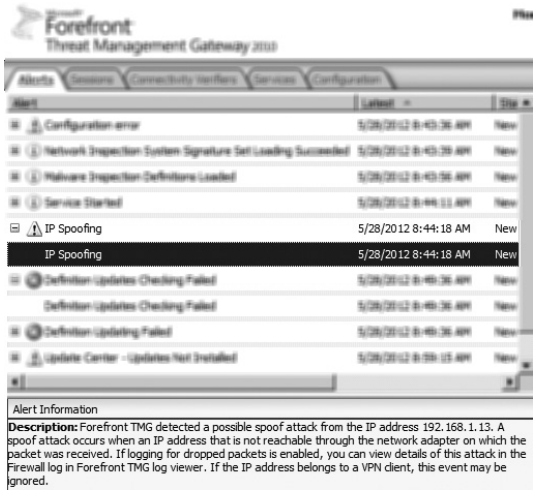

Used with permission of Microsoft Corporation

Figure 11-22 IP spoofing warning

6. In the left frame, click **Intrusion Prevention System**. In the Network Inspection System (NIS) tab, you can see some of the signatures that were loaded during the initial installation. Explore these signatures. In the middle frame, scroll to the right until the CVE Numbers column appears. Examine the Common Vulnerabilities and Exposures listings. In each item, the number after CVE is the year the listing was

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

released. Find one of the most recent listings, and double-click it. In the Signature Information Properties window, click the **Details** tab, and examine the information there. Copy the CVE number (for example, *CVE-2011-1977*), and use it as a search term on *cve.mitre.org*. Read the listing, and explore any links in the References section. In the Signature Information Properties window, click **Close**. Click the **Behavioral Intrusion Detection** tab.

7. Click **Configure Detection Settings for Common Network Attacks**. On the Common Attacks tab, notice that the enable detection box for Port scan is unchecked by default. Why do you suppose this is true? Click the **DNS Attacks** tab, and examine the content. In the Intrusion Detection window, click **Cancel**.
8. Click **Configure IP Options Filtering**, and examine the content of the IP Options tab and the IP Fragments tab. Click **Cancel**.
9. Click **Configure Flood Mitigation Settings**, and examine the content of the Flood Mitigation, IP Exceptions, and SIP Quotas tabs. Click **Cancel**.
10. In the left frame, click **Networking**, and examine the content on the Networks, Network Sets, Network Rules, Network Adapters, and Routing tabs.
11. In the left frame, click **System**. Click **Application Filters**, and examine the applications and protocols that are allowed access through TMG.
12. In the left frame, click **Logs & Reports**. With the Logging tab open, click **Start Query** in the right pane. Depending on what type of traffic is on your external network, you may see listings like those in Figure 11-23. Red listings indicate traffic that has been blocked. Black listings indicate allowed traffic. Click **Stop Query**.

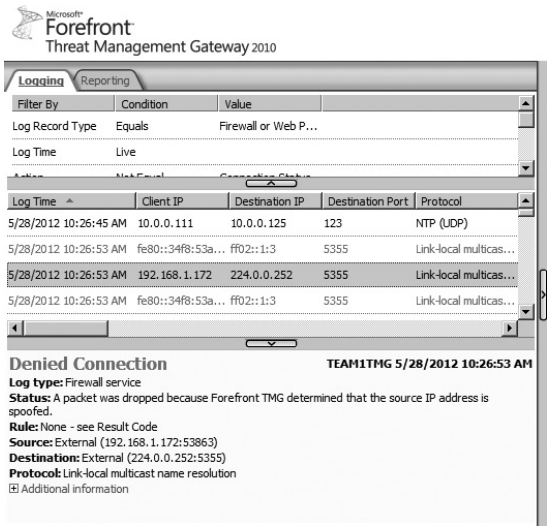

Used with permission of Microsoft Corporation

Figure 11-23 Firewall log

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

13. Click the **Reporting** tab. In the right frame, click **Create One-Time Report**. In the New One-Time Report Wizard window, type **Demonstration Report** in the Report Name text box, and click **Next**. In the Report Period window, verify that the Report period for this report is set to **Day**. In the Start date text box, enter the date you installed TMG. In the End date text box, enter today's date. Click **Next**. Accept the default selections in the Report Content window, and click **Next**. In the Send E-mail Notification window, click **Next**. In the Report Publishing window, type **\\Team*x*TMG\TMG_Reports** in the Published reports directory box. Click the **Publish using this account** box, and click **Set Account**. In the Set Account window, type **Team*x*\administrator** in the User text box. In both the Password and Confirm Password text boxes, type **Pa$$word**. Click **OK**. The Report Publishing window should look similar to Figure 11-24. Click **Next**. Click **Finish**. In the Forefront TMG window, click **Apply**, click **Apply**, and click **OK**.

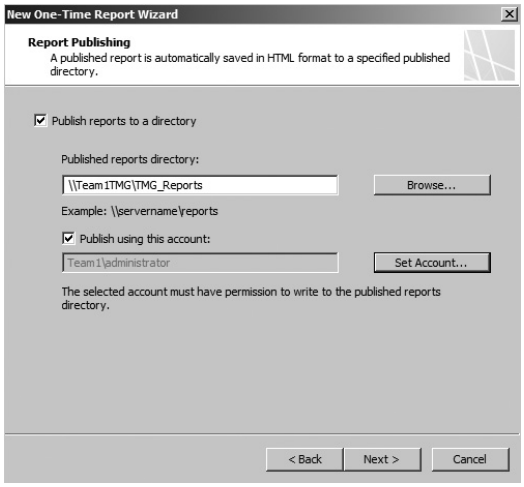

Used with permission of Microsoft Corporation

Figure 11-24 Report Publishing window

14. Leave the Forefront TMG window open. Click **Start**, click **Computer**, and double-click **Local Disk (C:)**. Right-click the white space in the right frame, click **New**, click **Folder**, and name the folder **TMG_Reports**. Right-click **TMG_Reports**, click **Properties**, click the **Sharing** tab, and click **Advanced Sharing**. Click the **Share this folder** box, click **Permissions**, and click **Add**. Type **Domain Admins** in the Enter the object names to select text box, click **Check Names**, and click **OK**. With Domain Admins selected, click the box under **Allow** in the Full Control row. Click **Everyone**, click **Remove**, click **OK**, click **OK**, and click **Close**. Close the Local Disk (C:) window.
15. Return to the Forefront TMG window. The Logs & Reports node should already be selected in the left frame, and the Reporting tab should be active in the middle frame. Right-click **Demonstration Report**, and then click **Generate and View Report**. The report opens in Internet Explorer. If you receive an error about ActiveX, click the yellow bar at the top of the browser, and then click **Allow Blocked Content**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Click **Yes** in the Security Warning window. Examine the report. Close the report after determining what data the report provides.

16. Click **Update Center**. Assuming that you have not had Internet access during these projects, you should see red alerts for Malware Inspection and Network Inspection System. Connect your system to the Internet. This may involve modifying the external interface IP address, default gateway, DNS server, and proxy settings. Check with your instructor for details.
17. In the right frame, click **Check for Definitions**. If the check is successful, the Definition Updates tab will look similar to Figure 11-25.

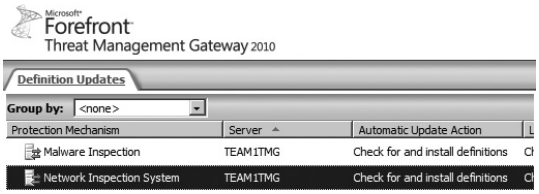

Used with permission of Microsoft Corporation

Figure 11-25 Definition updates completed

18. Log off all systems.

Case Projects

Case Project 11-1: VPN Filtering Rules

You work for a network consulting firm. Your client wants to implement a VPN system that supports PPTP-, SSTP-, and L2TP-based VPNs. The system will be used primarily for remote client access. You need to prepare a report for the client that explains the implications of providing this functionality in terms of firewall configuration. Your report must address the protocols and ports that need to be considered as the firewall policy is created. Prepare a one- to two-page report that addresses these issues.

Case Project 11-2: Threat Management Gateway Features

After receiving the report that you prepared in Case Project 11-1, your client has decided that an "all-in-one" product may be better than patching together a VPN infrastructure. You need to prepare a report that explains the features in the Microsoft Forefront TMG 2010 product. Your client has some technical understanding, but she needs to have the TMG features explained in a practical way so that she can consider how the product could assist in lowering costs and improving business processes. Prepare a two- to three-page report that meets the stated requirements.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 12

Internet and World Wide Web Security

After reading this chapter and completing the exercises, you will be able to:

- Describe weak points in the structure of the Internet
- Explain attack techniques against Web sites and Web users
- Explain methods for hardening Web and Internet resources

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

As Internet use becomes more prevalent, the number of attackers who take advantage of Internet users increases. In this chapter, you explore the infrastructure of Internet connectivity so that you can better understand weak points that attackers exploit.

You then explore several methods that attackers use against Internet components, such as Web servers, Web browsers, and e-mail clients. Finally, you learn how to harden infrastructure points to improve network security and reduce the risk of critical information and data being compromised.

Examining the Structure of the Internet

Internet use has increased exponentially in the past 10 to 15 years. Government organizations use it to streamline services and communications, for example, and businesses use it to conduct transactions and marketing operations. For many organizations, the Internet is indispensable in today's competitive marketplace.

Opportunists seek ways to exploit poorly designed systems on the Internet. Whether the aim is political, criminal, driven by greed, or just inspired by curiosity, the cycle of offense and defense is constant: Attackers discover new exploits, and then hardware and software vendors distribute notifications and patches to defend against these exploits. In the following sections, you examine the structure of the Internet and see how administrators and users can minimize risks.

Understanding the Structure of the Internet

The Internet is a group of networks tied together to form an infrastructure for communication. The terms Internet and World Wide Web (WWW) are often used interchangeably, but they are quite different. The Internet, a massive public medium established in the mid-1960s, is an interconnected web of networks and computers that work together to provide worldwide communication. The World Wide Web, which uses Hypertext Transfer Protocol (HTTP), is just one of the services the Internet offers. Many other services are offered, such as e-mail, which uses Simple Mail Transfer Protocol (SMTP) for communication, and file transfer, which uses File Transfer Protocol (FTP). The Web is a method of accessing information through the Internet by using HTML hyperlinks. It uses Web servers, Web browsers, and Web pages to communicate information through the Internet network.

The original Internet backbone was based on the U.S. Defense Department's ARPANET infrastructure and later the National Science Foundation's NSFNET, both of which relied on interconnected Internet backbones that support a distributed mesh topology. This early network defined the hierarchy of networks that is still in operation today.

By the early 1990s, the National Science Foundation (NSF) decided to stop funding NSFNET and move toward commercialization of the Internet. Many of the NSF's regional networks became commercial network service providers (NSPs), such as Netcom, UUNet/MCI Worldcom, and PSI Net. Also known as backbone Internet service providers, these NSPs expanded their backbone networks. To maintain their government funding, the NSF required the providers to allow free flow of traffic from one backbone to the other.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Tier System

The Internet is a tier system that connects networks around the world. This system starts with a backbone network connected via network access points (NAPs) to regional Internet service providers (ISPs). Regional ISPs service **point of presence (POP) ISPs** that connect to business, education, or home networks. The Internet is composed of a collection of backbones serviced by major carriers. The following sections explain NAPs and ISPs in more detail.

Routers and the Internet Communication Backbone The Internet communication backbone is an interconnected network of backbones owned by businesses or NSPs. Much like an anatomical backbone, where nerve signals travel across the spinal cord, an Internet backbone provides a conduit for network communication between different points of Internet access.

Routers direct network traffic to its destination via routing tables and updates from routing protocols. Routers in NSP backbones differ from routers in a LAN by the high amount of traffic they are designed to handle and the routing protocols they use for the Internet backbone environment. The physical memory, CPU speeds, interfaces, and operating systems (OSs) of routers used in NSP backbones can support enormous amounts of traffic and large routing tables.

Network Access Points **Network access points (NAPs)** are highly secure public facilities in which backbones have interconnected data lines and routers exchanging routing and traffic data. NAPs provide physical space, power, and network connectivity between different levels of the Internet's tier system, such as between a regional ISP and a POP ISP or between a backbone and a regional ISP.

Backbones exist in different regions of the world, and NAPs are positioned in each country to provide interconnectivity between these backbones. Each NSP backbone exchanges routing and traffic data in one of two ways: via NAPs or private peering relationships (see Figure 12-1). Private peering relationships are contracts between commercial NSPs or ISPs that enable them to bypass the Internet backbone for data and route exchanges.

Internet Service Providers An ISP provides access to the Internet at different levels depending on the type of ISP. A local ISP or POP ISP provides Internet access directly to consumers or businesses. A regional ISP sells bandwidth to local or POP ISPs or to organizations with high bandwidth requirements. A backbone ISP or NSP gives regional ISPs backbone access. If an ISP is large enough, it might offer services at all these levels, from local/POP access to backbone ISP.

Domain Name System **Domain Name System (DNS)** is a name-resolution service that translates fully qualified domain names, such as www.cengage.com, to IP addresses used to identify host computers, such as 69.32.133.79. Thirteen root servers named A through M are operated by commercial, educational, and government organizations to form the foundation of the Internet DNS. These servers are often targets of attack.

DNS is a hierarchical system, as shown in Figure 12-2. Root servers know which servers on the Internet are responsible for top-level domains (such as .org or .com). Each top-level domain has its own servers that delegate responsibility for domain name-to-IP address resolution to name

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

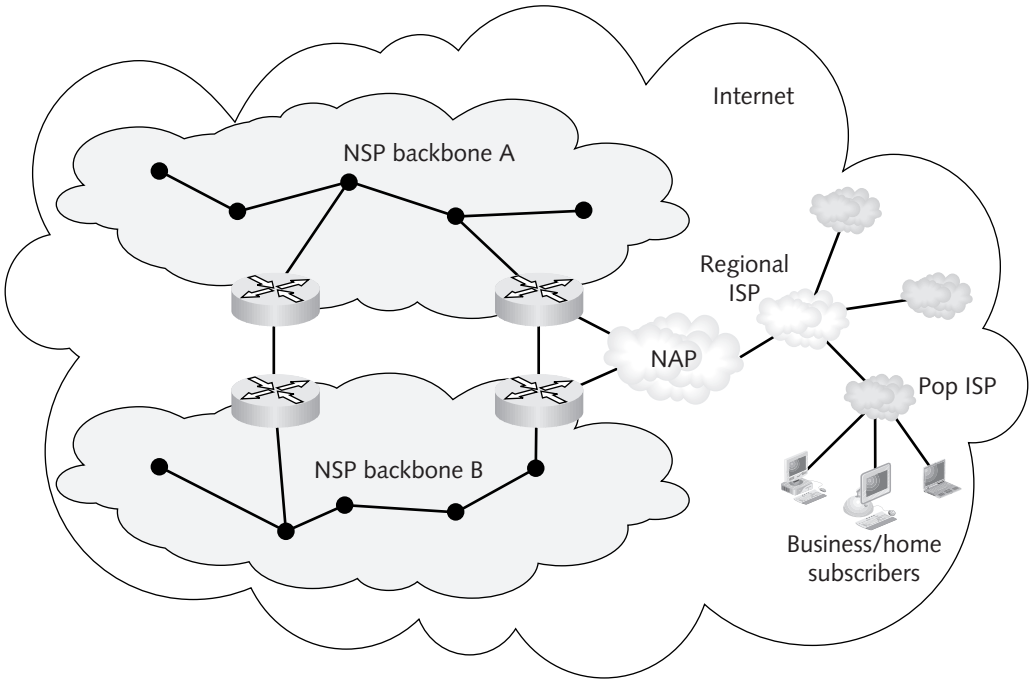

© Cengage Learning 2014

Figure 12-1 The Internet backbone: a network of NSP backbones

servers lower in the hierarchy. In reality, most DNS information on the Internet is cached, so DNS lookups rarely have to go all the way up the hierarchy to root servers.

The 13 root servers were located in the United States originally. However, servers C, F, I, J, K, and M are on other continents now, which is possible because they make use of **anycast addressing**. Unlike unicast and multicast addressing, anycast addressing enables any group of servers to act as a root server, regardless of location. Anycast addressing is a way of decentralizing DNS services, and balancing the load among several servers improves availability.

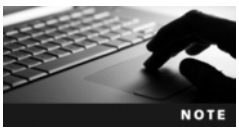

You can learn more about the DNS root server system at *www.root-servers.org*.

At ISPs, a local DNS server replicates entries from servers that are higher in the DNS hierarchy and resolves DNS requests. This server also forwards queries up the hierarchy if it cannot resolve a request with its own DNS table or cache.

Understanding Weak Points in the Internet's Structure

For all the usefulness the Internet offers, risks are inevitable when people and organizations operate in this environment. Attackers constantly discover new ways of exploiting the Internet infrastructure, and IT professionals must often play a catch-up game to stay ahead

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

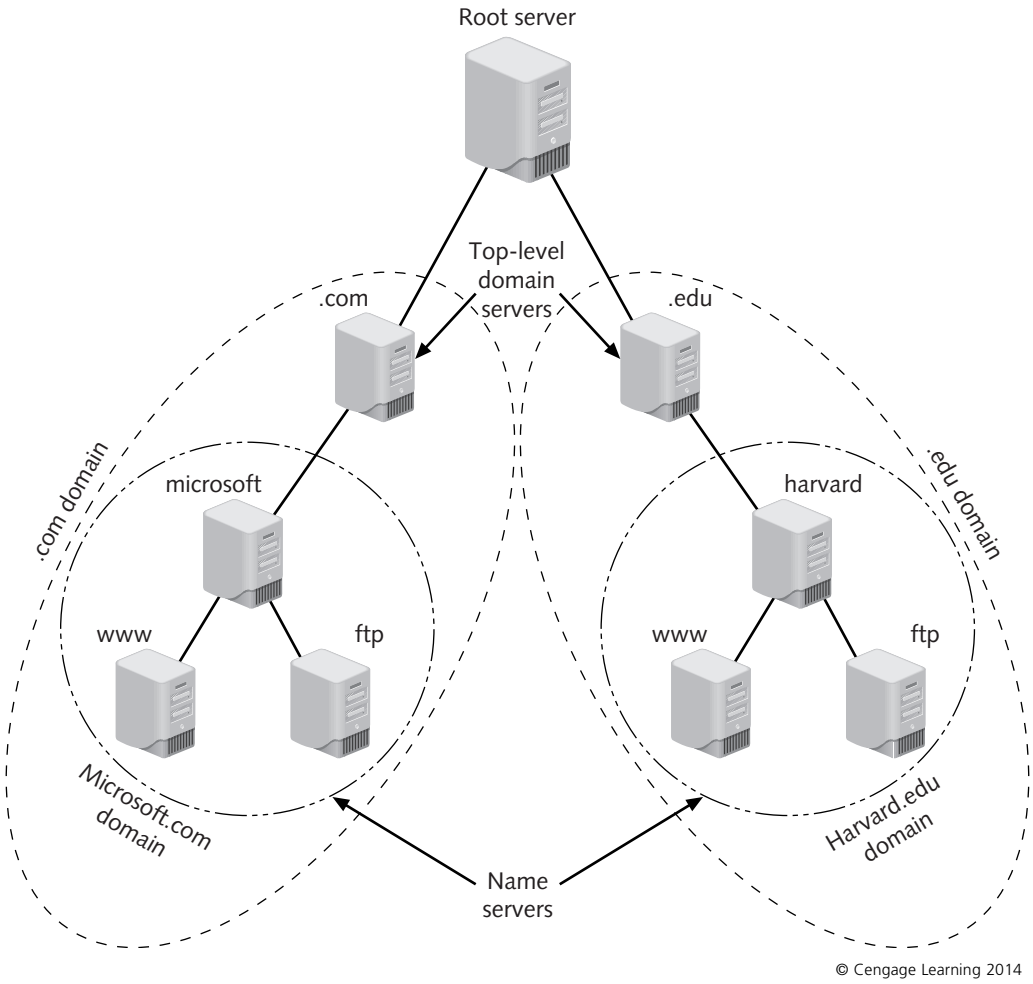

Figure 12-2 DNS hierarchy

of attackers' exploits. In the following sections, you examine some techniques that attackers use to exploit weaknesses in the Internet's structure.

IP Spoofing Computers on the Internet are identified mainly by their IP addresses, which are not authenticated by TCP/IP. This lack of authentication makes IP spoofing possible. Attackers change the source IP address in the headers of malicious packets they are sending to match a trusted host's IP address. To find a legitimate IP address, attackers sometimes send ping packets into a network and wait for responses. IP spoofing is most often used in denial of service (DoS) attacks. Attackers do not care about receiving responses to their packets; they want to trick a network's defense systems into accepting packets so that they can flood the network with packets and cause it to crash. Attackers sometimes use IP spoofing to fool servers into sending responses to their forged addresses, however. Because IP spoofing is so widespread, it makes accountability difficult for malicious actions on the Internet; someone who has intentionally accessed a restricted site can claim to be the victim of IP spoofing.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Modern network routers and firewalls offer software protection against known forms of IP spoofing. Packet filtering through routers is a major defense, for example. It includes ingress filtering to prevent spoofed IP source addresses from entering a network and perhaps egress filtering to prevent users inside a network from sending spoofed packets outside the network. In addition, IPv6 offers improvements to prevent IP spoofing, such as authenticated headers in each packet.

Routing Security The Internet network is linked by routers. Routing protocols such as Border Gateway Protocol (BGP) are used to communicate information updates for routing tables. However, routing information is not authenticated, so it is vulnerable to compromise. Attackers could send modified routing updates that misdirect data to a destination of their choosing. With this data, they can launch DoS attacks, use IP spoofing to intercept packets, or launch man-in-the-middle attacks, in which the attacker's computer or router is placed between the source and destination of a communication to intercept and steal information.

DNS Security DNS was originally designed as a public database for name-resolution services, so checking the authenticity and integrity of information stored in name servers was not considered necessary. This lack of security has caused several problems for the Internet community.

One problem is DNS cache poisoning, also known as DNS spoofing. When a name server is queried for DNS information that is not in its cache, it queries other servers. Because DNS information is not authenticated, attackers can send false data to a name server; the DNS cache is said to have been "poisoned." Attackers often use cache poisoning to steer unsuspecting victims to a server of their choice instead of the Web site where users intended to go.

A DNS name server for an organization contains database entries about every host on the network. Another problem caused by lack of authentication is DNS information leakage, which might occur if attackers access the database and use it to map target systems in the network. This information can be partially secured in DNS by blocking zone transfers, which replicate a name server's DNS information to other servers. Attackers could still retrieve the information by using DNS tools to query systems in an organization's IP namespace one by one until they capture a complete listing of DNS information. However, this method is slower and more painstaking.

Internet Host Security One of the bastions of Internet interconnectivity is the millions of host computers worldwide, but ironically, they are the weakest point of the Internet infrastructure. Attackers hijack many unprotected computers around the world and use them as "zombie" computers to deliver spam e-mail, DoS attacks, and malicious code. Attackers often assemble these zombies into **botnets** (networks of zombie computers) to magnify the scope and intensity of their attacks. According to M86 Security Labs, 91 percent of spam e-mail sent in May 2012 was delivered by hijacked zombie computers.

For the Internet and its Web component to function, many computers must be connected globally into one network. Each computer differs in the way it has been prepared to handle dangers on the Internet. For example, the risks that a careless Internet user takes in the United States, China, or South Africa can have dire consequences for careful users whose host computers are also connected to the Internet. The dangers of virus, Trojan, or DDoS

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

attacks on your system might exist simply because you are connected to the Internet. Therefore, good computing practices to minimize risks, such as antivirus software, firewalls, and system patches, are essential to withstand attacks from the Internet.

Web Site Attack Techniques

To exploit the Internet's weaknesses, attackers use a variety of innovative techniques. In the following sections, you examine how attack techniques are constructed and then learn how to best defend against them. First, you examine attacks against Web servers, and then you learn about attacks on the client side of the Internet: Web browsers and e-mail applications.

Attack Techniques Against Web Servers

The World Wide Web operates on the principles of a client/server network, and its basic building blocks are Web servers and client computers. Because millions of these network components are distributed around the world and their hardware and software configurations are so similar in nature, they are the Internet components that attackers target most often. Attackers probe common hardware/software server configurations, such as Windows running Internet Information Services or Linux running Apache Web Server, in an attempt to discover security holes.

Sensitive transactions, such as banking and e-commerce, are commonplace now on the Internet, so attackers often select Web servers that handle these transactions as targets for identity theft. The following sections explain ways that a Web server can be compromised.

Buffer Overflow Attacks

A buffer overflow attack exploits software vulnerabilities over which users and even network security personnel have little or no control. These common attacks often come with no warning and are almost impossible to detect and fix. These attacks have been used since the mid-1980s, when attackers discovered how to manipulate computer memory remotely by using worms and Trojans.

Generally, commercial OSs, Web servers, and databases are more vulnerable to buffer attacks than customized software that companies create for internal use. The source code is wrapped in a "black box" to protect it from tampering, but many attackers have the skill to access this code. News of pinhole vulnerabilities in black boxes travels quickly through the hacker community, and decompile tools for commercial software are readily available. After attackers have access to an application's code structure, they can look for weaknesses and errors in the source code—the root of buffer overflow attacks.

Keep in mind that buffer overflows themselves are not a cause for concern. These errors are common in computer operations. You have seen system errors similar to the one reported in Figure 12-3; they indicate that the buffer might have been corrupted during normal operations. The security problem starts when attackers discover poorly written code that causes buffer overflows and then inject malicious code into this breach.

A buffer is a section of random access memory shared by application processes that depend on one another but operate at different speeds or with different priorities. A buffer's purpose

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

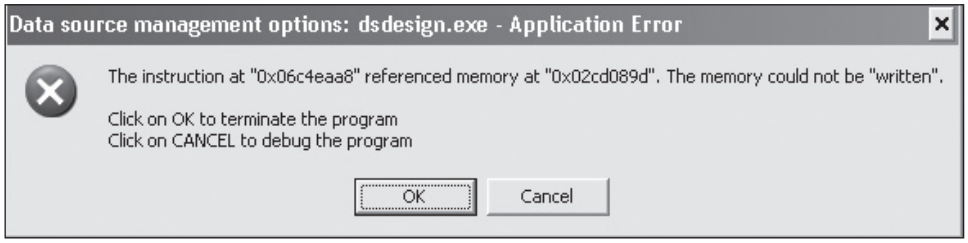

Used with permission from Microsoft Corporation

Figure 12-3 A common system memory error message

is to coordinate data intended for use by separate activities. A critical buffer component for this coordination is the call or function stack, and buffer overflow attacks are usually aimed at this component. Such attacks are sometimes called stack-smashing attacks.

The stack stores information about the processes an application is currently running as well as return addresses, local variables, and parameters. A return address is where a process goes when it finishes running. For example, if Process1 calls Process2 several times, Process1 pushes the instruction for the return addresses of Process2's calls onto the stack. When Process2 finishes running, it pulls its return address from the stack, and program control is transferred back to Process1's return address.

Stacks are allocated a fixed size in memory when they are created. Local variables for each running process occupy (or are "buffered in") this fixed memory space. If the process of pushing instructions onto the stack consumes all the space allocated for the stack, a buffer overflow occurs. The problem is that the program code does not adequately police how local variables are loaded and used in the buffer. For a buffer to be effective, its size and the algorithms for moving data in and out need to be protected, so good buffer design in programming is crucial in preventing attacks.

More specifically, buffer overflow attacks exploit a lack of "bounds checking" on the size of data stored in a buffer array. By writing data larger than the size allocated to an array, the attacker causes the buffer to overflow. For example, look at the following program code:

```
void buffSmash(void)
{
char *strOne = "abcdefghijklmnopqrstuvwxyz0123456789";
char buffOne[15];
strcpy(buffOne, strOne);
return;
}
```

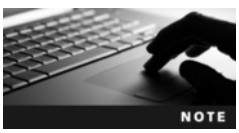

The code examples are shown in the C language because many C programs have buffer overflow vulnerabilities caused by a lack of bounds checking in some library functions, such as strcopy and gets.

The buffOne variable has been allocated 15 bytes in memory to store data assigned to it. The strOne string is larger than 15 bytes, but the code attempts to stuff it into buffOne's memory space. The result is a buffer overflow: buffOne spills over into memory space allocated for

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

another purpose. This error is only an annoyance until attackers discover it, and then it becomes a security problem.

In a buffer attack, intruders might be looking for an error message such as the one shown previously in Figure 12-3, which reveals a critical instruction pointer (0x06c4eaa8). Attackers attempt to overwrite this pointer with a return instruction injected through the buffer that redirects the process to run malicious code—with the ultimate goal of seizing control of the program and server.

Defending against a buffer attack is usually reactive rather than proactive because you cannot foresee vulnerabilities in commercial software. The best defense is to install patches and updates as soon as they are available. Most buffer attack damage is inflicted on unpatched systems. On the proactive side, installing intrusion detection and prevention software can be beneficial.

SQL Injection Attacks

Structured Query Language (SQL, pronounced *sequel*) is used to communicate with most relational database management systems (RDBMSs), such as Oracle, MySQL, SQL Server, and DB2. Because SQL is used so widely, particularly in e-commerce databases, it is a favorite target of attackers for data theft and destruction. As you have learned, buffer overflows are a result of poor coding. The same is true for SQL injection attacks: Web sites that have not been sanitized correctly are vulnerable to attack. The term *sanitized* describes computer applications or processes that have been protected against attacks.

A buffer overflow attack requires programming expertise, which limits the number of potential attackers. SQL injection, however, is plaintext scripting that is easy to learn and apply, making it a favorite language for fledgling attackers (often called script kiddies). No special tools are needed—just a computer connected to the Internet, a Web browser, and patience. The good news is that, unlike buffer code written by a third-party programmer, Web pages usually consist of custom-written code. Therefore, coding vulnerabilities can be fixed if they are detected in time.

SQL injection does not attack a Web server directly. It attacks the database used to support Web sites housed on the Web server, and more sophisticated attacks can be extended to attack the database server and its partner Web server. The next sections examine two SQL injection attack methods and their effects on database and Web servers.

SQL Injection: Web Form Attacks Web forms used to gather information, such as login pages or order forms, are potential entry points for attackers probing for Web site vulnerabilities. These forms are usually connected to a Web server's database, and a verification process checks information entered in the form and rejects incorrect entries. If the form's entry text boxes are not verified correctly, however, attackers can use them to send malicious code to the database, the database server, and perhaps even the partner Web server.

If attackers are targeting a specific Web site, any form on the site is a starting point to look for vulnerabilities. For example, attackers could start with a login form, e-mail form, or any information form. More likely, however, attackers have to search for vulnerable candidates, which requires patience. A common method of finding candidates is to use a Google search for login pages.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

To find login pages, go to *www.google.com* and try the following search terms: login.htm, login.html, login.asp, login.aspx, login.php, and login.jsp. These keywords return thousands of possibilities. You can also substitute *index* and *logon* for *login*.

The attacker might work through the Google results one by one or pick certain types, such as bank or e-commerce sites, and then start probing. The simplest method is to enter a single quotation mark (') in the first text box on a form and submit it. An unsanitized form allows the transmission of an unpaired single quotation mark in the SQL command to the database, which generates an error message. The database expects to receive a SQL command that is similar to the following:

```
SELECT somefields FROM sometable WHERE field1 = 'username' and field2 =
'password'
```

At this point, an attacker has no idea what the table or field names are, but is hoping to generate an error to find information. Here's an example of the SQL command the attacker sends to the database:

```
SELECT somefields FROM sometable WHERE field1 = ''' and field2=''
```

Notice the three single quotation marks following the equal sign after field1. They represent an unpaired quote, which generates an error message similar to that shown in Figure 12-4.

With this input, the attacker might be able to learn the following:

- The Web page is not well protected from intrusion.
- The database uses SQL Server and the Web server uses Internet Information Services.
- A careless administrator has not changed the default database username (sa).
- Pages are constructed with Active Server Pages (ASP), which could be a clue about the coding languages used on this Web site.

Another approach, or one that supplements the previous probe, is to inject SQL code that always evaluates to true (such as x=x). Attackers probe for familiar SQL patterns used to perform certain Web functions, such as verifying usernames and passwords or checking for valid e-mail addresses. Here is an example of an attacker's entry in a text box that requests an e-mail address:

SELECT somefields FROM sometable WHERE field = 'xyz@home.com' OR 'x=x'

It does not matter what e-mail address or text is entered between the first and second single quotation marks (in this case, xyz@home.com). The attacker is speculating about the SELECT statement's structure and has added another SQL command, `OR 'x=x'`. The statement x=x always evaluates to true, so this SQL statement asks for fields that either contain the bogus e-mail address or in which x=x. If either condition is met, a return should be made.

This query might generate different responses, such as delivering all rows of a table or displaying a message that a correct password has been delivered to a specified e-mail

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

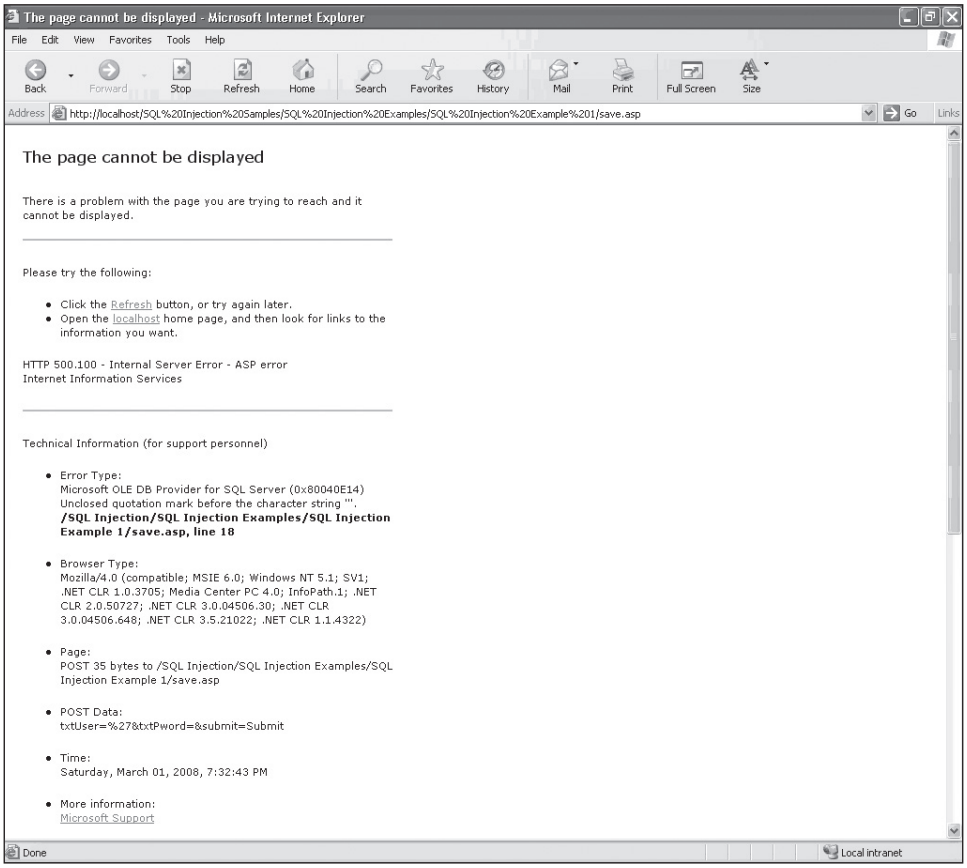

Used with permission from Microsoft Corporation

Figure 12-4 A database-generated error message

address. With either response, the attacker has gained more information for creating a database footprint—a map of tables and fields in the database—that can be used to launch more serious exploits.

This method requires patience, but attackers with destructive intent could learn enough to cause serious damage. For example, if attackers know the type and name of the database and table names, they might be able to enter a SQL command to shut down the database or drop a table, as shown in this example:

```
SELECT knownfield FROM knowntable WHERE knownfield='' or 1=1; DROP TABLE
knowntable; --' AND otherfield=''
```

The semicolons denote the start of a new SQL command, and the double dash (--) indicates that the following text is a comment, not code.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

With this access, attackers could add users to tables or print lists of information, such as credit card numbers. The most serious threat is an attack that gains administrative access to a database and stored procedures used to issue server commands, which could provide access to the entire Web site operation.

SQL Injection: Query String Attacks The second method used for SQL injection attacks involves the query string used to send information to a database. When a user clicks a link on a Web page, information is sent to the Web server. For instance, on a retail site, a user might click a product picture to see more information, which is stored on a database that supports the Web server. The product's ID code and perhaps other information is attached to the Web page address and sent to the Web server for action. The information being sent is clearly visible in the browser address bar and can be the source of an attack on the Web site's database.

For example, the Web address of an item at an e-commerce site could be www.commercesite.com/php?id=55167. This address would request the Web server to deliver information from a table row with a unique ID of 55167. An attacker could attempt to manipulate the underlying database by modifying the ID number and sending the request again. In a secure system, an error such as "No records matched your query" would result. The error would reveal no information about the Web server software or the database configuration. However, in a less secure system, an attacker may be able to learn important information that helps develop an attack.

Attackers use this method to probe Web databases for vulnerabilities; it's the same technique used in Web form attacks with a different injection point. To generate more information from an error message, attackers might try to add extra code to the end of the query string, as shown:

```
index.asp?id=1 AND password=1
```

As with the Web form attack, the goal of any query string probe is to gain additional information about a database's structure for further attacks.

Defenses Against SQL Injection Attacks SQL injection attacks are isolated to custom applications, so administrators can prevent them, unlike buffer overflows, which require third-party vendors to make code adjustments. The first course of action is to prevent malicious code from being entered in Web pages that allow user input. A common mistake is for site administrators to stop there and take no further protective measures. Attackers can exploit SQL in many other ways, however, so you should take the following steps to close all potential holes:

- Tighten database authentication and limit table access. Always require password access to the database, and never leave default usernames set up during installation. Most attackers are familiar with the default administrative username sa, so make sure to change it.
- Use stored procedures to eliminate passing any SQL commands to the database.
- Validate all user entries to make sure they are formed properly. Perform this validation in several places if necessary. There should be two layers of validation: form-level validation at the browser before the Web page is submitted and server-level validation when the information reaches the server for processing.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Place the Web server and database server in a network DMZ.
- Use nonstandard naming conventions in database construction. To thwart attackers, you should make database names, table names, and field names difficult to guess.
- Inevitably, database errors do occur, so configure a custom error message that does not reveal information for attackers to exploit. The standard 404 error message often reveals server information that attackers can use.

With these simple precautions, you can immunize a Web server and its database server against SQL injection attacks.

Attack Techniques Against Web Users

In the previous section, you learned about attack techniques that are directed against computer or network systems. Now you examine attacks directed at Web users through commonly used applications, such as Web browsers and e-mail programs. These attacks fall into the category of social engineering because they prey on emotions such as curiosity, anxiety, fear, and greed. Unlike some attacks against systems, almost all attacks against Web users can be prevented.

Every user takes on a measure of security risk when interacting on the Web. Attacks on Web users center on a variety of objectives, including identity theft and simple malicious behavior. Informed Web users should understand these attack methods and know how to prevent them. In the following sections, you explore some of these risks and examine measures that can eliminate or at least minimize the risks.

Phishing Attacks

Phishing is an attack through a Web browser that displays false information masquerading as legitimate data. Phishing is a deception designed to steal personal information such as credit card data, account numbers, usernames, and passwords.

Phishing attacks can take many forms that range from simple to quite sophisticated. A simple form that has persisted for many years is the Nigerian money scam. The perpetrator sends millions of e-mails to a random selection of addressees, asking for help in transferring a large sum of money from Nigeria to the United States. To be "rewarded," the e-mail recipient must provide personal banking information to assist in making the transfer. The scam has grown in popularity to include letters from China, North Korea, and Russia. The letters are almost comical in their wording, and it is hard to imagine that anyone would take them seriously. People are still taken in by this scam, however.

Another form of phishing involves Web page deception. An attacker sends millions of e-mails that appear to come from trusted sources, such as banks, insurance companies, or payment sites like PayPal or eBay. The phony e-mails can look legitimate and include corporate graphics that mimic those of a real organization. Figure 12-5 shows a typical corporate phishing e-mail.

These e-mails have the following characteristics:

- The e-mail is unsolicited and unexpected.
- The logo and other graphics are copies of corporate images and seem to be legitimate.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

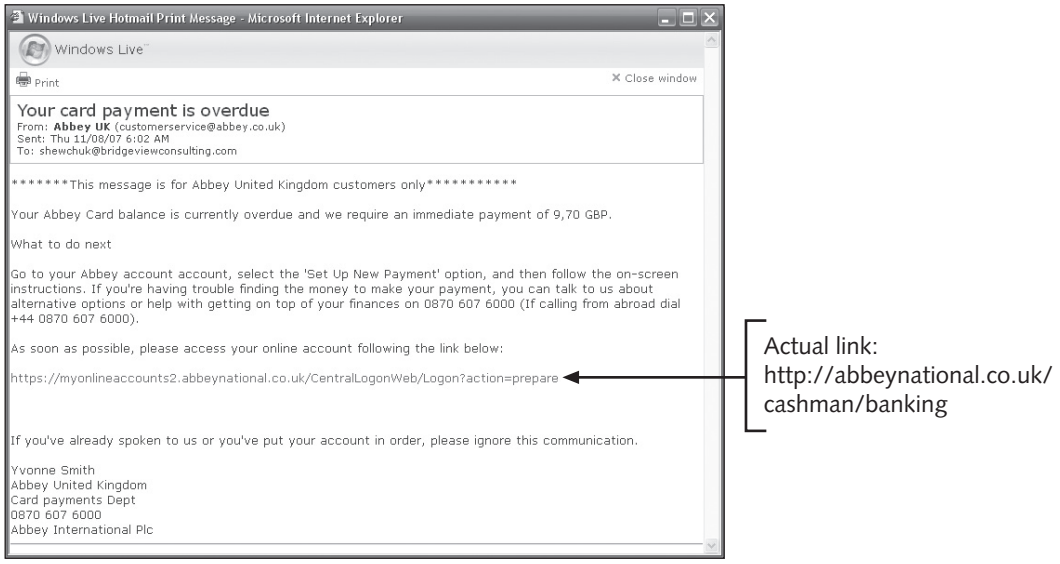

Used with permission from Microsoft Corporation

Figure 12-5 A typical corporate phishing e-mail

- The message uses a generic greeting, such as "Dear valued customer" or "Corporate bank user," instead of the recipient's real name.
- The message conveys a sense of urgency, such as "Please respond immediately," "Your account has been locked," "Your credit card has expired," or "There is a security problem." The intent is to make readers hurry and perhaps overlook illegitimate aspects of the e-mail.
- Personal account information is requested, usually by asking that the information be confirmed.
- The e-mail contains a link that seems to be a secure HTTPS link. When you navigate to a Web page with this link, a lock symbol depicting an SSL-encrypted site appears at the lower right. If you hover the mouse pointer over the link, however, the real Web page address appears at the lower left. Other links in the e-mail, such as "Security Policy" or "For further information," point to the same illegitimate link.
- Usually the link to which you are redirected is no longer active after several hours. Phishing attackers play a game of hit-and-run to elude authorities, so they do not remain stationary for long.

In the early days of phishing, e-mail messages were filled with grammatical mistakes and spelling errors, and the graphics were poor copies. Attackers have become more sophisticated, however, and now their e-mails are often well crafted and difficult to discern as fakes.

The objective of a phishing attack is to entice e-mail recipients to click the bogus link, visit the fake Web site, and enter personal information that can be used to steal personal assets or identities. These attacks are prevalent because they are easy and inexpensive to create and carry out. All an attacker needs is an e-mail list and a Web site address, which can be purchased for less than $10 or even stolen by hacking into a legitimate Web site. E-mail

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

lists are collected in several ways, such as surfing through newsgroups or blogs or purchasing lists from spammers or other phishers. Some members of newsgroups or other Internet discussion forums make their e-mail address more difficult to harvest by spelling out the address as follows: "dean underscore farwood at comcast dot net."

A couple of phishing variations have surfaced. In **pharming**, traffic to a legitimate Web site is redirected to the attacker's Web server. To do this, the attacker modifies the user's host file, a local file on the hard drive that assigns an IP address to a domain name. The attacker can also use DNS manipulation, such as cache poisoning. Pharming is more insidious than phishing because it does not require users to click a URL to go to the attacker's server.

Another variation, **spear phishing**, focuses on specific targets, usually financial institutions or government agencies. The attacker identifies users or groups in an organization by using common avenues, such as e-mail, telephone, Facebook, and corporate Web pages. He might even gather names and job titles from assigned parking spaces. The attacker then mounts a campaign to exploit employees' vulnerabilities with the goal of using their access rights to penetrate the corporate infrastructure. Even the best corporate security measures are vulnerable to this attack because it is an inside job: It is launched within the corporate security perimeter by using employees' authentication credentials. Spear phishing is becoming more common because people are increasingly careless about revealing personal information on chat sites, blogs, and retail shopping sites.

Preventing most phishing attacks is a matter of educating Web users. Train employees to follow these simple guidelines for preventing phishing attacks:

- Check the browser address bar and footer. If a Web site does not have an HTTPS address or a lock icon, it is not a secure site, and no personal information should be entered. Also, the lock icon can be faked, so users should not believe that its presence indicates a secure site.
- If you get an e-mail from a company you are familiar with, call the company to confirm that the request is legitimate. Companies usually have an 800 number for this purpose.
- Forward any obvious phishing e-mails to the company being portrayed in the phishing attempt. PayPal and eBay, for example, have forwarding addresses set up for this purpose.
- Delete any unsolicited e-mails about foreign banking.

As a network security administrator, you must frequently remind users of these simple preventive measures. Send regular e-mail reminders so that prevention against phishing attacks becomes second nature.

To see actual phishing e-mails, perform a Google search using the company name and the word *phishing*.

File Attachment Attacks E-mail attachments are a common vehicle for introducing malicious code into a network. These attacks first occurred in 2002 when JPEG attachments were discovered to have virus code embedded in the file header code. Until then, data files had been relatively immune to infection.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The attack requires two virus components. The first part spreads in the form of a traditional Win32 executable virus, arriving via e-mail or portable media. This virus makes changes to the Registry so that JPEG files are run through an extractor before they are displayed. The virus strikes if the user attempts to view a JPEG image and the extractor finds the second virus component in the graphics file header. Having both parts on the same computer is rare, however, because standard virus protection typically detects the presence of a virus arriving via e-mail or portable media. However, users should be cautious about viewing image file attachments from unknown sources.

A more recent file attachment attack victimized the security company RSA. An e-mail with the subject line "2011 Recruitment Plan" was sent to RSA employees. Attached to the e-mail was a file called 2011 Recruitment Plan.xls, which contained an embedded Flash file that installed a backdoor program. Using this backdoor, the attackers were able to install remote access tools, capture user credentials, and eventually exploit a system that contained data about RSA's SecureIT product. The backdoor enabled successful attacks on a number of U.S. military contractors' networks.

ActiveX Control Attacks An ActiveX control is a Windows object coded in languages such as C++, Visual Basic, and Java. Its purpose is to deliver dynamic, interactive content to Web pages. A control object is compiled and stored in a CAB file, which is stored on a Web server and accessed by referencing the object's assigned CLASSID, as shown in the following code from the Developer Connection QuickTime site:

```
<OBJECT CLASSID="clsid:02BF25D5-8C17-4B23-BC80-D3488ABDDC6B" WIDTH="160"
HEIGHT="144"
CODEBASE="http://www.apple.com/qtactivex/qtplugin.cab">
<PARAM name="SRC" VALUE="sample.mov">
<PARAM name="AUTOPLAY" VALUE="true">
<PARAM name="CONTROLLER" VALUE="false">
<EMBED SRC="sample. mov" WIDTH="160" HEIGHT="144" AUTOPLAY="true"
CONTROLLER="false"
PLUGINSPAGE="http://www.apple.com/quicktime/download/">
</EMBED>
</OBJECT>
```

When used for legitimate purposes, an ActiveX control can be a beneficial addition to a Web site; however, attackers have discovered that an ActiveX control can be programmed to run malicious code on a user's Web browser. ActiveX controls do not require user action to be activated. They run automatically when the browser loads the Web page that contains them. ActiveX controls have almost full access to the Windows OS and can perform many functions, including running code on an unprotected computer, which could involve accessing and downloading files, planting Trojan programs and worms, or destroying system programs.

The defense against malicious ActiveX controls is to make sure that you scrutinize them by using security settings on Web browsers. Browsers can be set to block ActiveX controls from running on Web pages, for example. You can also adjust browser settings to permit certain types of ActiveX controls to run and block others.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Java Applet Attacks A Java applet is a small program sometimes used as embedded code in Web pages. Java applets were considered immune to hacking because they were encased in a "sandbox" that could not interact with a host computer outside the confines of a browser; they could communicate only back to their code base. However, Java applets have been used with Internet Explorer and Netscape to exploit the OS and access system files. In the Internet Explorer attacks, malicious code embedded in a Java applet was used to exploit a proxy server network connection. The user's session was then redirected to a location of the attacker's choice without the user's knowledge, and the attacker was able to capture the user's information.

In the Netscape attacks, vulnerabilities in Netscape Communicator and Navigator made it possible for Java applet code to gain unauthorized local and remote file access. A malicious Java applet could read files from the local file system by opening a connection to a URL and gaining complete file access. The applet could then send files back to the server from which it originated.

Using this method, the connection is reversed, and the user of the browser sends information to the applet originator. This communication reversal negates the protection of the user's firewall, which watches for incoming vulnerabilities, not outgoing ones.

Although the combination of circumstances in these attacks is rare, it does emphasize the need to patch your system with the most recent updates and hot fixes. Furthermore, only signed applets should be permitted to run on Web browsers. In the Java applet exploits discussed in this section, software patches from vendors fixed the problem.

Hardening Web and Internet Resources

Establishing and maintaining a hardened network with secure hosts requires continuous vigilance and regular updates of components. New versions of software, hardware, and network media are released frequently, but the threats against networks and systems change just as often. A network security administrator alone cannot keep up with the daily deluge of security concerns. By enlisting the help of security experts and adopting a preventive stance toward network security, the task becomes far less daunting.

Seeking the assistance of security experts need not be an expensive venture. Most of the help you need is free. Your first stop should be the supplier of your firewall and antivirus software. All reputable vendors maintain informative Web sites with excellent guidelines for how to best use their products. They also offer automatic, timely downloads of the latest virus signature databases so that your systems are protected as soon as possible. With your network systems protected, you can then push the updates to all connected host computers automatically.

Hardening DNS Servers

When attackers probe networks to look for vulnerabilities, they pay special attention to servers that host Internet services, such as DNS or Web servers, because they store valuable personal and corporate information. This section explains techniques for hardening these servers against attacks.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

A primary DNS server is authoritative for specific domains and has DNS zone files that change as needed. A zone file is a set of instructions for resolving domain names into IP addresses. A secondary DNS server receives a read-only copy of the zone file to improve the query performance of DNS services. An internal DNS zone file contains entries of all internal hosts on a network, and an external zone file contains only host entries that are visible to the public. A **zone transfer** occurs when a zone file is sent from the primary DNS server to secondary DNS servers for updating. A zone is just an abbreviated way of referring to the domain name for which a DNS server is configured. If the domain name is myschoolsite.edu, for example, its components are the name (myschoolsite) and the generic top-level domain (edu). A subdomain might be staff.myschoolsite.edu, for example. Figure 12-6 shows an example of a small zone file.

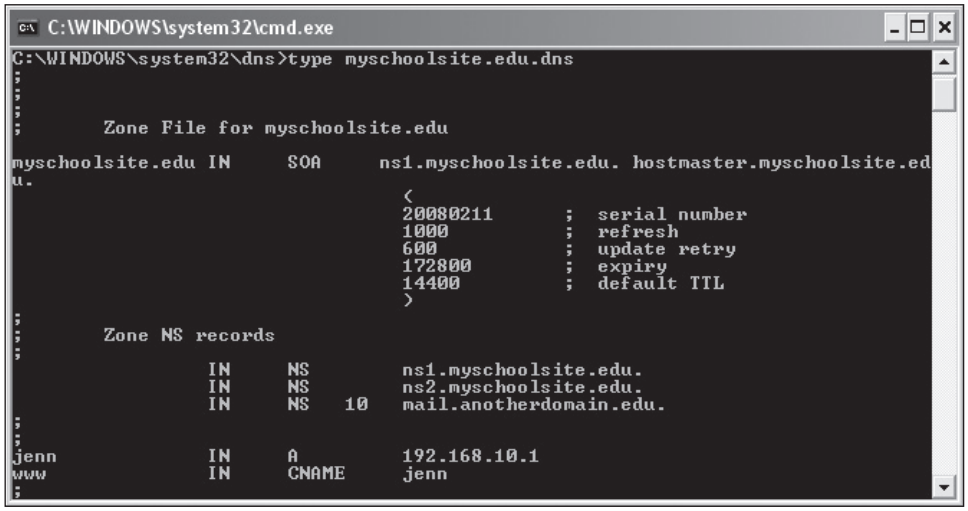

Used with permission from Microsoft Corporation

Figure 12-6 A zone file for myschoolsite.edu

If zone transfers are not secured, attackers might be able to intercept them and retrieve a complete listing of network resources and possible targets for attack. One of the most serious mistakes a network administrator can make is to allow untrusted Internet users to perform zone transfers. Transfers should be allowed only between primary and secondary DNS servers. If the DNS server does not use a segregation method to separate external DNS information from private internal information, internal IP address and host name information could be exposed to attackers, who would then have an electronic road map of the organization. The following example shows probing code that an attacker might try as a starting point:

```
C:\Nslookup
```

This code might yield the following information:

```
Server: somedomain.com
Address: 10.10.10.10
```

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

With this information, the attacker could attempt to change the server to the network's primary DNS server and then list and pull records from it.

Securing zone transfers is usually straightforward; you simply configure all DNS servers to restrict zone transfers to specific authorized servers. Using selective transfers minimizes the risk of unauthorized users getting a copy of a zone file.

If your organization has a DNS server that is authoritative for your domain on the Internet, make sure that DNS servers are in a DMZ and that a split DNS architecture is used. A **split DNS architecture** physically separates public DNS servers from the organization's internal DNS servers. Public DNS servers are used for authoritative DNS services to the Internet. In addition, internal DNS servers use a non-Internet domain name, such as .corp or .local; the authoritative DNS servers conform to Internet domain requirements. A split DNS architecture like the example in Figure 12-7 prevents internal zone information from being stored on an Internet-accessible server and prevents internal DNS entries from being sent over Internet DNS.

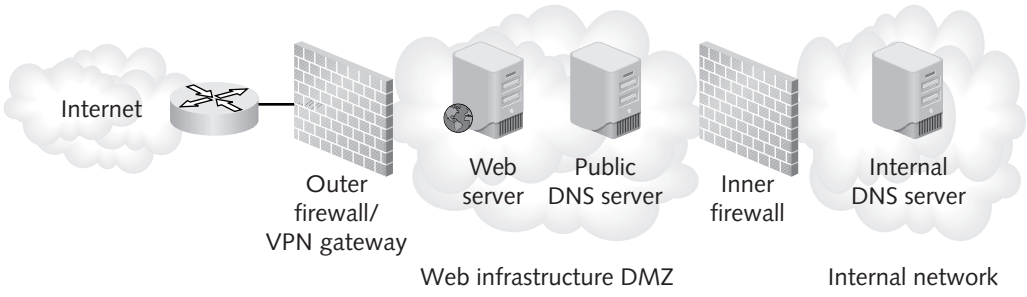

Figure 12-7 A split DNS architecture

A variation on this split DNS architecture is called **split brain DNS architecture**. In this system, a physical separation still exists between internal and external DNS servers, but both DNS systems use the same domain. While this system creates less confusion for internal users because they do not have to distinguish between internal and external DNS domains, it puts added responsibility on the DNS administration team, who must make sure that all references to internal hosts are removed from the DNS zone maintained in the DMZ.

Another layer of DNS security is at the network perimeter, where you should configure firewalls and routers with rules or filters to prevent zone transfers to the Internet. Zone transfers typically occur over TCP port 53.

DNSSEC

The DNS protocol was designed long before the designers of the Internet became concerned about security. Not surprisingly, the protocol has been found to be vulnerable to exploitations, as you learned earlier in this chapter. These exploitations can result in loss of confidentiality (tapping Voice over IP), Web site impersonation, e-mail hijacking, DNS cache poisoning, and theft of information, including logon credentials and credit card data. **DNS Security Extensions** (**DNSSEC**) was created to thwart some of these attacks. The goals of

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

DNSSEC are to provide authentication of DNS data, ensure integrity of DNS data, and authenticate the denial of existence of DNS data. It is important to note that DNSSEC addresses only these goals; it does not provide message confidentiality or protect against DDoS attacks.

DNSSEC uses cryptographic techniques to provide security for DNS data. Digital signatures are created and stored in a new type of DNS resource record—an RRSIG record—as part of the process of creating a signed DNS zone. A **security-aware resolver** is a system that is compliant with DNSSEC and that attempts to use a DNS server to resolve a fully qualified domain name to an IP address (or vice versa); this system can access the RRSIG for the DNS zone in question. To authenticate the address resolution, the resolver can access the public key associated with the RRSIG record by accessing the DNSKEY resource record where the public key is stored. A Public-key Infrastructure (PKI) exists to provide a chain of authentication for these keys. Just as a system that uses typical asymmetric encryption must be configured with at least one root key, a security-aware resolver must be configured with at least one DNSSEC **trust anchor**. A trust anchor is the top-level digital certificate in the PKI chain.

Larger ISPs have begun the process of implementing DNSSEC. This development is important because most DNS queries are serviced by ISP-maintained DNS resolvers. For example, Comcast, a major ISP in the United States, began migrating customers to DNSSEC resolvers in late 2010. By January 2012, Comcast had created signed DNS domains for all the domains under its control and had migrated all customers to DNSSEC-validating resolvers.

As you learned earlier, DNSSEC does not provide data confidentiality. Another weakness is that an attacker may be able to enumerate the contents of a DNS zone by following the NSEC resource record chain. An NSEC resource record is the Next Secure record that allows a resolver to trace the authentication path of the RRSIG. Also, DNSSEC is considerably more complicated than traditional DNS, which increases the possibility of errors and therefore a loss of service. Finally, the effectiveness of DNSSEC depends on unbroken chains of authentication. Until all Internet DNS zones are DNSSEC compliant, there is no assurance that the goals of DNSSEC can be met.

Hardening Windows Web Servers

Regardless of its platform, a Web server is usually secured by hardening the underlying OS, installing patches, disabling unused services, and restricting the number of user accounts and their access permissions. In addition, you can use platform-specific software tools.

Internet Information Services (IIS) is the Web server used in Windows 2000, Windows XP Professional, Windows Server 2003 and 2008, Windows Vista, and Windows 7.

Authentication Authentication is an important consideration in all forms of information security, but considering that Web servers are often accessible by untrusted users, it is particularly important when configuring Web server security. IIS 7 allows you to select one of two forms of authentication:

- Challenge-based authentication, in which the Web client must respond to a challenge from the Web server. An example is Integrated Windows Authentication, in which Active Directory credentials are used.
- Login redirection-based authentication, in which users must enter credentials on a login page. IIS 7 also supports SSL-based digital certificate authentication.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Many Web servers do not formally authenticate users; they permit anonymous authentication. E-commerce sites, for example, want users to access their Web sites without authentication until the users intend to buy merchandise. In IIS 7, these users are logged in with the account IUSR.

Windows Basic Authentication requires that users enter a username and password, but it is not browser specific. The downside is that it transmits passwords in plain text. Windows Digest Authentication uses Active Directory to authenticate users, but the client browser must support the HTTP 1.1 protocol. Windows Authentication supports both Kerberos and NTLM (New Technology LAN Manager, a legacy authentication method). Extended Protection is an authentication method available in IIS 7.5. Extended Protection is designed to decrease the risks associated with man-in-the-middle attacks by providing additional information, such as channel-binding tokens and service-binding identifiers.

Access Control IIS 7 allows you to restrict access to the Web server based on IP address, IP address ranges, and domain names. Access can be limited based on other parameters as well, such as computers, groups of computers, or domains. Also, access is granular, meaning it can be restricted to certain Web sites, applications, directories, and individual files. Furthermore, you can filter the types of HTTP requests that will be processed by the Web server.

Data Confidentiality IIS supports SSL encryption, so you can request and install Internet server digital certificates, install domain server digital certificates, or even create a self-signed server certificate. IIS configuration of SSL includes the ability to require SSL to access the Web server, determine the bit length of cryptographic keys in use, and control whether clients must verify their identity to connect with the Web server.

Controlling Dynamic Content Windows Web servers have traditionally used Internet Server Application Programming Interface (ISAPI) and Common Gateway Interface (CGI) to provide interactive and dynamic content. ISAPI extensions are applications, and ISAPI filters are programs used to modify or enhance IIS functionality. Both of these ISAPI components are typically implemented as .dll files. CGI is a standard commonly used in the creation of interactive forms; it is implemented as .exe files. IIS 7 allows you to restrict the activity of ISAPI and CGI components.

Shared Configuration Today, relatively few single Web servers are found in e-commerce settings. Server arrays and large server farms are common. IIS 7 supports shared configuration, which allows administrators to import configuration files and cryptographic keys from a centralized location. Similarly, configuration files and keys can be exported from a single server to the central location as a backup.

Other Security Considerations IIS security features are important to consider when planning Web server security, but they do not guarantee a hardened system with the fewest possible vulnerabilities. You should also follow these precautions:

- The underlying Windows OS must be hardened and maintained by installing the latest service packs, patches, and hot fixes, and by removing or disabling unnecessary services.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- A domain controller should not also function as an IIS Web server. Domain controllers store Active Directory information and control network access and authentication. This server handles critical services, so making it available on the Internet is not a good security practice. Domain controllers should be kept in the protected internal network and separated from the Internet with firewalls.
- Place the Web server in a secure room. An organization cannot achieve information security without physical security. Despite all the passwords and firewalls used on a network, a Web server can still be compromised if it is kept in an unsecured area. Restrict access to Web servers by using physical security measures, such as surveillance cameras, locks, smartcards, and other access control systems.
- Do not connect the IIS Web server to the Internet before it is fully hardened. When a server is connected to the public Internet, it usually does not take long for an attack to occur.
- Remove NTFS write and execute permissions when possible to minimize the risk of unauthorized users changing files or running programs.
- Grant permissions for modifying and viewing IIS logs to system and local administrators only. This precaution makes it harder for attackers to modify log files to hide their activities. As an added precaution, store logs on another server, not the IIS Web server.
- Allow only the administrator to log on locally to the Web server. Secure services outside the OS, such as SQL Server, to prevent them from being exploited as user accounts.
- If you are serving Web pages to the Internet, place the Web server in a firewall-protected DMZ.

As with all Microsoft products, service packs, patches, and hot fixes for IIS are released periodically. Installing them as soon as possible is important, especially when they address security issues. Subscription to the automatic Microsoft Security Notification Service is recommended for updated news on IIS as well as updates on all Microsoft products.

Configuring Security Settings in Apache Web Server

Apache Web Server, the most widely used Web server application, is installed mainly on UNIX and Linux systems, although a Windows version is available. Apache's vulnerabilities are not publicized as much as those for Windows Web servers, but they do exist. Some misguided Apache administrators believe that Apache is secure out of the box and do not pay much attention to hardening, which can be a grave mistake. As with any Web server, Apache requires hardening to ensure security for Web sites and users. The Center for Internet Security (CIS) recommends the following security settings for Apache:

- Harden the underlying OS as you would any OS by removing unused applications and sample code and updating OS patches and hot fixes.
- Install the latest Apache binary distribution code from the OS vendor. This approach is usually easier than compiling your own binary code for installation because the vendor has already done most of the configuration work for you.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Disable unnecessary Apache modules and services, disable processing of server-side includes (SSIs), and delete unneeded or default Apache files and sample code. These measures reduce the number of Web processes that are available to attackers.
- Create Web groups so that users can be granted limited administrative rights without having root access.
- Create user and group accounts with limited privileges for running Apache Web Server, and never run Apache as the root account. If the Web service runs with root permissions, any compromise results in attackers having root access to the Web server.
- Subscribe to OS vendor and Apache security advisories to stay informed about security issues.
- Develop customized messages for Web pages that display error information. As you learned previously, attackers can use error messages to gather information about server setup.
- Install the ModSecurity module to have URLs in Web traffic inspected for anomalies. For example, an attacker might send the following URL request to a Web server to delete the accounts database table:

  http://www.myweb.com/login.asp?username=admin';DROP %20TABLE%20accounts--

  The ModSecurity module adds a filter to prevent these types of requests.
- To secure access, use Digest authentication instead of Basic authentication for accepting usernames and passwords. (If you need to review these authentication methods, conduct an Internet search for articles.)
- When setting access control lists (ACLs), determine whether allow or deny rules are evaluated first. An ACL's effect could change if you do not use the correct order of evaluation.
- Use Secure Sockets Layer (SSL) to encrypt the communication from user to Web server. First, download and install the mod_ssl module from *www.modssl.org/source/*. Then install an SSL certificate purchased from a recognized certification authority (CA), such as VeriSign or Geotrust. The CA includes directions for installing the certificate. For example, VeriSign provides instructions at *https://knowledge.verisign.com/support/ssl-certificates-support/index?page=content&id=AR212*.
- Limit the Web server to accepting and processing only certain HTTP request methods, such as GET, POST, HEAD, and PUT.
- Disable HTTP traces to prevent attackers from investigating HTTP request paths for potential targets. An HTTP trace asks the Web server to echo back an HTTP request's contents and is often used for debugging. Attackers could use this information to access sensitive data, such as authentication data or cookies from an established connection. For more information, see *https://www.kb.cert.org/vuls/id/867593*.
- Enable logging on the Web server so that you can spot potential problems and suspicious activity. If the server is compromised, logs also give you a record for forensics analysis. To prevent attackers from accessing and altering logs, store them on a separate network server, not the Web server.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Chapter Summary

- The Internet is an interconnected web of networks and computers that work together to provide worldwide communication. In this tiered system, the Internet backbone is connected via network access points (NAPs) to regional ISPs, which service point of presence (POP) ISPs that connect to business, education, or home networks.
- Domain Name System (DNS) is a hierarchical system that provides name-resolution services for translating host names to IP addresses. The foundation of the DNS hierarchy is 13 root servers located around the world. Root servers communicate with servers lower in the hierarchy that are responsible for top-level domains; these servers delegate responsibility for domain name-to-IP address resolution to name servers further down the hierarchy.
- Internet weak points are caused by problems with IP address authentication, routing protocol security, DNS security, and Internet host security.
- TCP/IP does not authenticate IP addresses, so attackers can change the source IP address in headers of malicious packets to match a trusted host's IP address. This practice is called IP spoofing. In addition, routing information is not authenticated, so it is vulnerable to compromise.
- DNS was originally designed as a public database for name-resolution services, so checking the authenticity and integrity of information stored in name servers wasn't considered necessary. This lack of security has resulted in DNS cache poisoning attacks and DNS information leakage.
- The millions of host computers around the world are the weakest point of the Internet infrastructure. Attackers hijack unprotected computers and use them as zombies for delivering spam e-mail, DoS attacks, and malicious code. These zombies are often assembled into botnets (networks of zombie computers) to magnify the scope and intensity of attacks.
- Web servers are the Internet components that attackers target most often. Attackers take advantage of poor programming practices and unpatched system vulnerabilities to target Web servers.
- A buffer overflow attack exploits coding flaws in common commercial software, such as OSs. Attackers try to use this flaw to generate system memory errors in a running process and then inject malicious code through the memory buffer.
- A SQL injection attack uses plaintext scripting in an effort to generate information that attackers can use to destroy data, disrupt Web site operations, and launch further attacks. Because this attack targets custom-written code, vulnerabilities can be fixed if they are detected in time.
- There are two types of SQL injection attacks. In Web form attacks, malicious SQL code is passed to servers through login pages or entry text boxes on Web page forms. In query string attacks, malicious SQL code is added to query strings (URLs) in a Web browser's address bar.
- Web user attacks exploit social engineering techniques to target users and take advantage of vulnerabilities in Web browsers by using e-mail, ActiveX controls, Java applets, and file attachments.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Phishing is an attack through a Web browser. Phishing uses e-mail to send users false information masquerading as legitimate data in an attempt to steal personal information such as credit card data, account numbers, usernames, and passwords. Two variations of phishing are pharming and spear phishing.
- ActiveX controls do not require user action to be activated. They run automatically when the browser loads the Web page that contains them, and they have almost full access to the Windows OS. Therefore, attackers can program ActiveX controls to run malicious code on Web browsers, plant Trojan programs and worms, or destroy system programs.
- Attackers have been able to use embedded code in Java applets to exploit the OS and access system files. These attacks involve redirecting user sessions in an effort to capture user information.
- To harden DNS servers, make sure you allow zone transfers only between primary and secondary DNS servers. In addition, place DNS servers in a DMZ and use a split DNS architecture to physically separate public DNS servers from internal DNS servers. At the network perimeter, configure firewalls and routers to prevent zone transfers to the Internet so that DNS information is not revealed to attackers. DNSSEC is a method of authenticating and ensuring the integrity of DNS data, but it has not been fully implemented yet.
- IIS 7 has several features that allow you to improve security, including controls for authentication, encryption, authorization, and access. You should take additional preventive actions, such as hardening the underlying OS, placing the server in a secure room, and positioning the server in a firewall-protected DMZ.
- The CIS recommendations are helpful guidelines for configuring server processes to harden Apache Web servers. You can also install the ModSecurity module to add firewall and intrusion-detection capabilities.

Key Terms

anycast addressing A network addressing scheme that allows DNS services to be decentralized among a group of servers, regardless of their location.

botnets A network of zombie computers that attackers assemble to magnify the effect of an attack.

Domain Name System (DNS) A hierarchical name-resolution service for translating host names to IP addresses; used mainly on the Internet.

DNS Security Extensions (DNSSEC) A protocol designed to improve DNS security by using cryptography to ensure DNS integrity and authentication.

network access points (NAPs) Highly secure public facilities where commercial Internet backbones and ISPs exchange routing and traffic data.

pharming A variation of phishing that intercepts traffic to a legitimate Web site and redirects it to a phony replica site.

phishing Using social engineering techniques via e-mail to trick users into providing personal information at the attacker's Web site. This site is designed to look like a legitimate business site.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

point of presence (POP) ISPs ISP facilities that provide connectivity to the Internet for business, education, and home users.

security-aware resolver A computing system that is compliant with DNSSEC and that attempts to resolve a fully qualified domain name to an IP address (or vice versa).

spear phishing A variation of phishing directed at specific users instead of using spam e-mail. Spear phishing is often directed at employees of a particular organization, for example.

split DNS architecture A network architecture that divides DNS services between two servers: a public DNS domain with a server on the organization's DMZ for Internet services and an internal DNS domain with a server on the internal network for service to internal hosts.

split brain DNS architecture A network architecture that uses a single DNS domain with a DNS server on the organization's DMZ for Internet services and a DNS server on the internal network for service to internal hosts. All records that refer to internal hosts must be removed from the DMZ DNS server.

trust anchor The top-level digital certificate in a PKI chain.

zone transfer The communication of a zone file from the primary DNS server to secondary DNS servers for updating.

Review Questions

1. The Internet backbone is connected to regional ISPs via which of the following?
   a. POP ISPs
   b. network service points
   c. network access points
   d. carrier network points

2. How many root servers are in the DNS infrastructure?
   a. 10
   b. 11
   c. 13
   d. 14

3. Attackers can exploit routing information updates to do which of the following? (Choose all that apply.)
   a. Launch DoS attacks.
   b. Poison DNS caches.
   c. Use IP spoofing to intercept packets.
   d. Launch man-in-the-middle attacks.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

4. Attackers often use DNS cache poisoning to do which of the following?
   a. Query systems on a network one by one.
   b. Steer unsuspecting users to a server of their choice instead of the Web site where users intended to go.
   c. Flood the network with packets and cause it to crash.
   d. Install a virus on the network.
5. Which of the following is caused by a flaw in how a running process allocates memory to a variable?
   a. unsecured cryptographic storage
   b. buffer overflow
   c. broken authentication
   d. SQL injection
6. Which of the following is a common type of SQL injection attack? (Choose all that apply.)
   a. Web form attack
   b. browser executable attack
   c. system tray attack
   d. query string attack
7. In a SQL injection attack, which character is an attacker most likely to use?
   a. asterisk
   b. single quotation mark
   c. exclamation mark
   d. double quotation mark
8. Which of the following attack methods target Web users? (Choose all that apply.)
   a. social engineering
   b. phishing
   c. SQL injection
   d. pharming
9. What is a requirement for a successful file attachment attack?
   a. The user must open the file attachment.
   b. The user must reply to the e-mail that contains the attachment.
   c. The user must delete the file attachment immediately.
   d. The attachment must be an image file.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

10. Which of the following factors enables attackers to program ActiveX controls to run malicious code on a user's Web browser? (Choose all that apply.)
    a. ActiveX controls run in a sandbox that allows interaction with the OS.
    b. ActiveX controls do not require user action to be activated.
    c. ActiveX controls run automatically when the browser loads the Web page that contains them.
    d. ActiveX controls have almost full access to the Windows OS.
11. A Web server can be hardened just by configuring the Web application correctly. True or False?
12. For optimum efficiency, configure a domain controller to function also as an IIS Web server. True or False?
13. When securing an Apache Web server, which of the following tasks is not necessary?
    a. installing the latest Apache patches
    b. disabling processing of server-side includes (SSIs)
    c. deleting unneeded or default Apache files and sample code
    d. creating a privileged user ID for the Apache Web User account with root access
14. In a DNS zone transfer, what is actually transferred?
    a. fully qualified domain names and IP addresses
    b. usernames and passwords
    c. server MAC addresses
    d. UDP and ICMP messages
15. To keep log files organized, store them on the server you are monitoring. True or False?

Hands-On Projects

Hands-On Project 12-1: Finding Domain Information

Time Required: 10 minutes

Objective: Find your network's DNS and ISP information.

Description: In this activity, you use network tools from *Network-Tools.com* to discover domain information about your network.

1. Log on either to the Windows Server 2008 or Windows 7 system.
2. If necessary, configure your network interface for Internet access.
3. Start your Web browser and go to *http://network-tools.com*. Figure 12-8 shows the interface for this tool.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

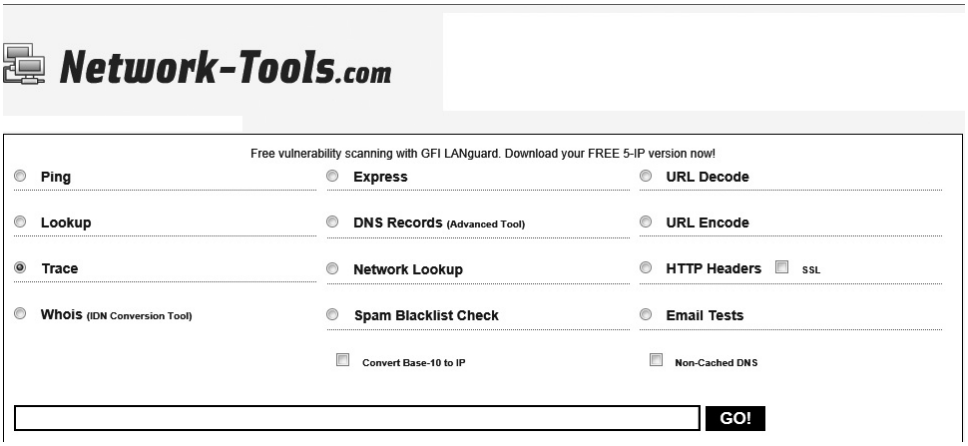

Source: Network-Tools.com

Figure 12-8 Network-Tools

4. Click the **Lookup** option button, enter your school's domain name (such as schoolname.edu) in the text box below the option buttons, and then click **GO**. What is the IP address that corresponds to your school's domain name? What country and region are shown for this domain name?
5. Click the **Whois** option button, and then click **GO**. Who is the contact for your school's domain name? What name servers are listed for your school's domain?
6. Click the **DNS Records** option button, and then click **GO**. What are the mail servers for your school's domain (MX records)?
7. Click the **Trace** option button, and then click **GO**. Which ISPs handle the school's Internet traffic?
8. Exit your browser, and leave your system running for the next project.

Hands-On Project 12-2: Examining Internet Explorer Security Settings

Time Required: 20 minutes

Objective: Become familiar with the security settings in Internet Explorer version 9.

Description: In this activity, you examine Internet Explorer's settings for handling security functions, such as browsing history, cookie management, security zones, and pop-ups.

1. If necessary, log on to Windows 7.
2. Start Internet Explorer. Verify the version number by clicking the **Tools** icon and then clicking **About Internet Explorer**. These controls are indicated by an "A" and "B" in Figure 12-9. If necessary, use Windows Update to upgrade Internet Explorer to version 9.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

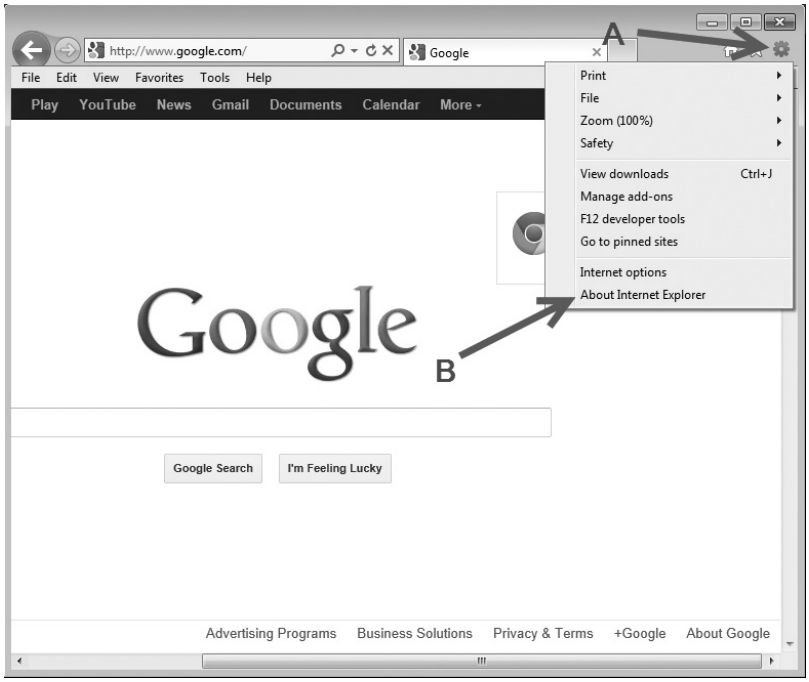

Source: Google; also, used with permission from Microsoft Corporation

Figure 12-9 Accessing the Internet Explorer version number

3. Click the **Tools** icon, and click **Internet options**.
4. Click the **General** tab, and examine the Browsing history section. Click **Delete**. You can delete temporary Internet files, cookies, and the history of the Web sites you have visited. In addition, you can control whether to save your Web site passwords, and you can determine whether to keep track of Web sites for which you want to allow ActiveX controls. Click **Cancel**.
5. Click the **Security** tab. Verify that the Internet icon is selected. In the Security level for this zone section, note that the default level for the Internet zone is Medium-high. What note is displayed at this level about the treatment of ActiveX controls? Slide the level controller up to High and down to Medium to observe the explanations of these levels. Note also that the Enable Protected Mode box is selected. This feature isolates Internet Explorer from system resources that could be exploited by malicious code.
6. Click **Custom level**. Here you can make specific changes to a variety of functions by disabling them, enabling them, or requiring them to request permission to run (prompt).
7. Scroll to the ActiveX controls and plug-ins section. For the first ActiveX control setting, Allow ActiveX Filtering, click the **Disable** option button. For all remaining ActiveX control settings, click the **Enable** option button. When you enable particularly unsecure settings, the color of the setting changes to alert you that the setting is not recommended.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

8. Click **OK** to close the Security Settings window. When you see a warning message about changing security settings for this zone, click **Yes**. Note that the Security level for this zone is now listed as Custom, and several indications show that the Internet zone setting puts your computer at risk. See Figure 12-10.

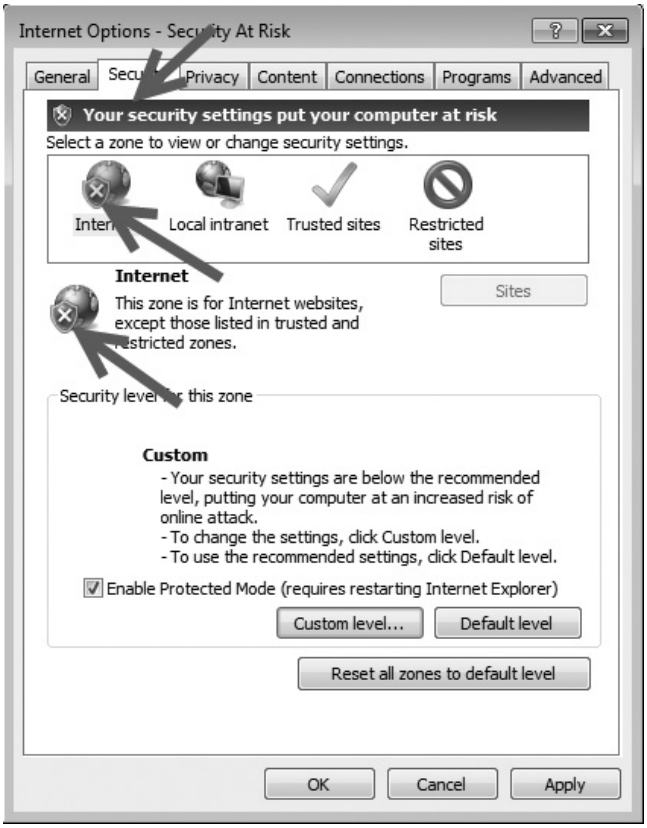

Used with permission from Microsoft Corporation

Figure 12-10 Internet Explorer Internet zone security warnings

9. Click **Local intranet** and examine the default settings for sites that you classify as being on a more trusted network: your company's intranet.
10. Click **Trusted sites**. Here you can configure the specific Web sites that you trust and that can have a lower security filtering level. Click **Sites**. Here you can add specific Web site addresses. Click **Close**.
11. Click the **Privacy** tab. In the Settings section, explore the various settings for the Internet zone by using the sliding control. These settings are mostly associated with cookies, the files placed on your hard drive by Web sites that you visit. You can also control whether Web sites can request access to information about your physical location and whether to allow or block pop-ups.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

12. Click the **Security** tab, click **Reset all zones to default level**, and click **Apply**. Note that the security warnings associated with the Internet zone are removed. Close the Internet Options window and Internet Explorer.
13. Leave your system running for the next project.

Hands-On Project 12-3: Installing Internet Information Services (IIS)

Time Required: 15 minutes

Objective: Install IIS.

Description: In this activity, you install Microsoft's IIS Web service and verify functionality of the default Web site created by the installation.

1. Log on to Windows Server 2008.
2. Start **Server Manager**, click **Roles**, and click **Add Roles**. In the Before You Begin window, click **Next**.
3. In the Select Server Roles window, click the **Web Server (IIS)** box, and click **Next**. Read the Web Server (IIS) window, and click **Next**.
4. In the Application Development section of the Select Role Services window, click **CGI**, **ISAPI Extensions**, and **ISAPI Filters**. In the Security section, check **Basic Authentication**, **Windows Authentication**, **Digest Authentication**, **URL Authorization**, and **IP and Domain Restrictions**. Click **Next**.
5. Click **Install**. When the installation is complete, click **Close**, and close Server Manager.
6. Click **Start**, click **Administrative Tools**, and click **Internet Information Services (IIS) Manager**.
7. Expand **TEAMXSRV**, expand **Sites**, and click **Default Web Site**. The screen should look like Figure 12-11. This is the configuration page for the default Web site created by the installation of IIS. Leave this window open.
8. Open **Internet Explorer**. In the address box, type **http://localhost** and press **Enter**. The default Web site's default home page should appear, as shown in Figure 12-12. Close Internet Explorer.
9. Leave your system running for the next project.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

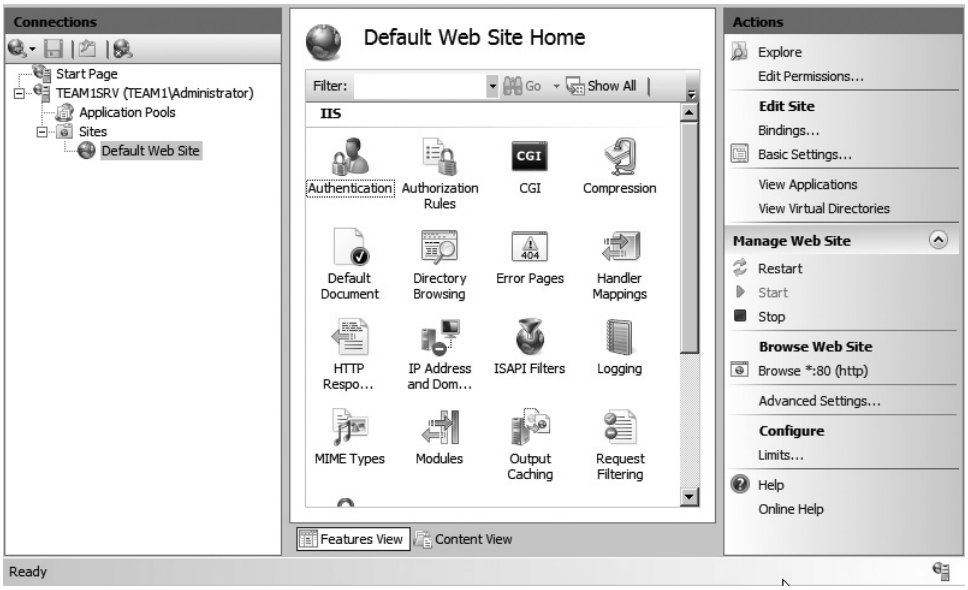

Used with permission from Microsoft Corporation

Figure 12-11 IIS default Web site home

Used with permission from Microsoft Corporation

Figure 12-12 IIS default Web site home page

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Hands-On Project 12-4: Configuring Security for Internet Information Services

Time Required: 60 minutes

Objective: Become familiar with security configuration in IIS.

Description: In this activity, you explore many of the configuration settings in IIS. You learned about some of these settings in this chapter. In particular, you implement authentication of Web site users and discover the difference between authenticating domain users and non-domain (public) users. Also, you configure SSL so that users must encrypt all data sent to the Web server and decrypt the data that is sent from the Web server. You configure the Web server to provide a self-signed digital certificate to support encryption. Normally, you would create a certificate that has a root certificate signed by a commercial certificate authority. That way, public users would already trust your Web server certificate. In this lab, you will see the difference between accessing a trusted Web site and an untrusted one.

1. Log on to your Windows Server 2008 system, if necessary.
2. Make sure that your network interface is on the same subnet as the Windows 7 system and that both the Windows Server 2008 and Windows 7 interfaces list the Windows Server 2008 IP address as the DNS server.
3. You should have the IIS Manager console open with the default Web site node selected. If not, click **Start**, click **Administrative Tools**, and click **Internet Information Services (IIS) Manager**. Expand **TEAM*X*SRV**, expand **Sites**, and click **Default Web Site**.
4. In the right frame, click **Bindings**. Select the **http** row and click **Edit**. Note that all unassigned IP addresses for this system are listening for HTTP requests at port 80. In the Host name box, type **TEAM*x*SRV.TEAM*X*.NET**, click **OK**, and click **Close**.
5. In the center frame, double-click **Authentication**. Select **Anonymous Authentication**. In the right frame, click **Edit**. Note that users who log on anonymously will be logged as IUSR. Click **Cancel**.
6. Note that anonymous authentication is the only method of access to the default Web site. Select **Basic Authentication** and then click **Edit** to determine what configuration is needed when enabling this type of authentication. Click **Cancel**. Repeat this process for Digest Authentication.
7. Select **Windows Authentication**, and click **Enable** in the right frame. In the right frame, click **Advanced Settings** and note the available options. Click **Cancel**.
8. Click **Anonymous Authentication**. In the right frame, click **Disable**. Click **Digest Authentication**, and click **Enable** in the right frame.
9. Open Internet Explorer, type **http://team*x*srv.team*x*.net** in the address box, and press **Enter**. You should be successful because you have already authenticated to the operating system with an Active Directory account and Windows Authentication is enabled.
10. Log on to your Windows 7 system as an administrator of the domain. Click **Start**, click **Control Panel**, click **Administrative Tools**, double-click **Computer Management**, and click **Local Users and Groups**. In the middle frame, double-click the **Users** folder. From the **Action** menu, click **New User**. In the New User window,

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

type **publicuser** in the User name text box. In the Password and Confirm password text boxes, type **Pa$$word**. Uncheck **User must change password at next logon**, click **Create**, and click **Close**.

11. Click **Start**, click the **Shut down** arrow, and click **Switch user**. Click **Ctrl+Alt+Del**, click the box above **Other User**, and log on with the username **teamxclient\publicuser** and the password **Pa$$word**. You have now logged on to Windows 7 with a local account, not one that will be recognized by Active Directory when you access the IIS server.
12. Start Internet Explorer. In the address box, type **http://teamxsrv.teamx.net**, and press **Enter**. Because anonymous access is not permitted, you must authenticate with an account in the Team*x* domain. Type **administrator** as the username and **Pa$$word** as the password.
13. If a window appears that contains the **Use recommended security and compatibility settings** option button, click the button, and click **OK**. A second tab opens to describe the Internet Explorer version 9 upgrade. If prompted to do so, enable add-ons. Close Internet Explorer. Log off as the public user. Leave Windows 7 running.
14. Access Windows Server 2008.
15. In the IIS Manager console, click **Default Web Site** in the left frame.
16. Double-click **Authorization Rules**. In the right frame, click **Add Allow Rule**. Here you can specify users and groups to which you want to allow access to Web content. Click **Cancel**. Note the Add Deny Rule link that allows you to specify users and groups to whom you want to deny access. Click **Default Web Site** in the left frame.
17. Double-click **CGI** in the middle frame, click **Help** in the right frame, and read about the CGI function. Close the Help window. Click **Default Web Site** in the left frame.
18. Double-click **IP Address and Domain Restrictions** in the middle frame. In the right frame, click **Add Deny Entry**. Here you can specify an IP address, IP address range, or domain names to restrict from accessing Web content. Click **Cancel**. Click **Default Web Site** in the left frame.
19. Double-click **ISAPI Filters** in the middle frame, click **Help** in the right frame, and read about the ISAPI filter function. Close the Help window. Click **Default Web Site** in the left frame.
20. Double-click **Request Filtering** in the middle frame. Just under "Use this feature to configure filtering rules" is a series of tabs. Use the right arrow at the end of the row to view all the tabs. Click **Help** in the right frame, and read all the Help pages under the Request Filtering Page section. Close the Help window. Click **TEAM*X*SRV** in the left frame.
21. In the middle frame, double-click **Server Certificates**. In the right frame, click **Create Self-Signed Certificate**. In the next window, type **Team*x* Web Server Certificate** to specify a friendly name, and click **OK**. In the right frame, click **View** and examine the certificate you just created. How many bits are used for the public key in this digital certificate? Click **OK**. Click **Default Web Site** in the left frame.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

22. Double-click **SSL Settings**. Note that because HTTPS is not bound to the server's network interface, SSL cannot be required. Click **Default Web Site** in the left frame. In the right frame, click **Bindings**. Click **Add**, and select **https** from the Type menu. Note that the port automatically changes to the HTTPS default port of 443. In the SSL certificate menu, select **Team*x* Web Server Certificate**. Click **OK** and click **Close**.
23. Return to **SSL Settings** and click the **Require SSL** box. In the right frame, click **Apply**.
24. Return to the Windows 7 system, and log on as publicuser. Start Internet Explorer. In the address box, type **http://team*x*srv.team*x*.net**, and press **Enter**. You should receive an error message similar to the one shown in Figure 12-13.

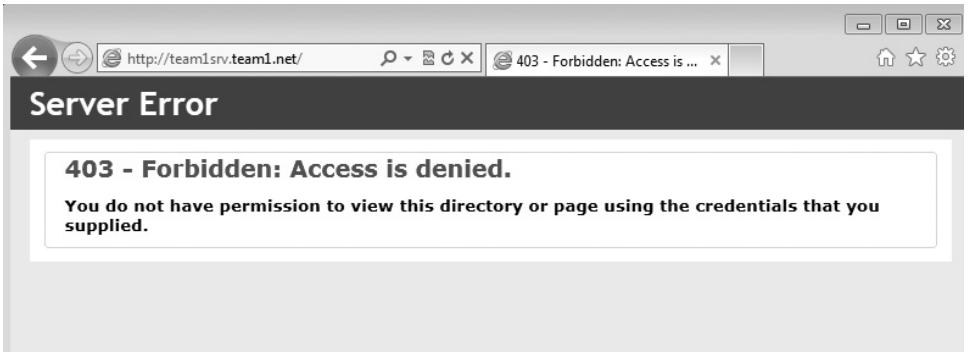

Used with permission from Microsoft Corporation

Figure 12-13 Access denied message

25. Attempt to connect to the Web server using the required HTTPS. Type **https://team*x* srv.team*x*.net** in the address box of Internet Explorer, and press **Enter**.
26. You receive a warning that you do not have a certificate that can verify the Team*x* Web server uses a trusted certificate. Click **Continue to this website (not recommended)**. In the Windows Security box, authenticate the username **administrator** and the password **Pa$$word**. This time you can access the Web site, but the address box has turned red as a warning that the Web site is not trusted. Click the red **X** in the address bar to read the detailed warning. See Figure 12-14. Close Internet Explorer and log off Windows 7. Keep Windows 7 running.
27. Return to Windows Server 2008.
28. Create a directory in the root of C: called **Certs**. Share this folder and give all users full control.
29. From the IIS Manager console, click **TEAM*X*SRV** in the left frame, and double-click **Server Certificates** in the middle frame.
30. Select the Team1 Web Server Certificate in the middle frame, and click **Export** in the right frame. Click the **...** button beside the Export to box. Browse to the **C:\Certs** directory, type **Team*x* Web Server Certificate** in the File name box, and click **Open**. In the Password and Confirm password text boxes, type **Pa$$word** and click **OK**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

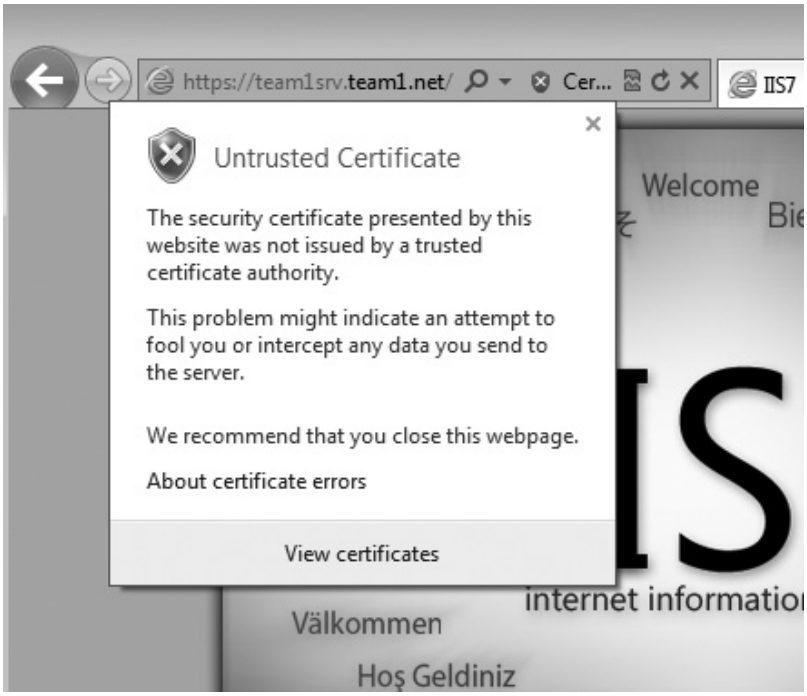

Used with permission from Microsoft Corporation

Figure 12-14 Untrusted Certificate warning

31. Return to Windows 7 and log on as publicuser. Click **Start**. In the Search programs and files text box, type **\\Team*x*srv\Certs** and press **Enter**. When the authentication window appears, log on as **administrator** with the password **Pa$$word**. You should see the Team1 Web Server Certificate on Windows Server 2008 in C:\Certs.
32. Right-click **Team1 Web Server Certificate**, and click **Install PFX**.
33. In the welcome window, click **Next**. Note that the remote certificate file is listed in the File name text box. Click **Next**. In the Password text box, type **Pa$$word**, and click **Next**.
34. In the Certificate Store window, click the **Place all certificates in the following store** option button, click **Browse**, click **Trusted Root Certification Authorities**, and click **OK**. Click **Next**, and click **Finish**.
35. Read the Security Warning window, and then click **Yes**. Click **OK** when informed that the import was successful.
36. Start Internet Explorer, type **https://team*x*srv.team*x*.net** in the address box, and press **Enter**. This time, no warning appears about an untrusted certificate because you have imported the server certificate and now trust it. Authenticate as the administrator. Users now are required to use SSL encryption when connecting to your Web server.
37. Log off all systems.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Case Projects

Case Project 12-1: Web Server Security Analysis

A challenge that faces the information security community is to make sure their organizations' top decision makers understand the importance of effective security policies. Understandably, managers who are responsible for an organization's financial stability want to ensure that investments in "nonrevenue" activities are necessary and effective. In many ways, the field of information security is in its infancy; practitioners do not have a long history of research that indicates the most effective security countermeasures against specific threats. Also, the threats are constantly changing. As new generations of information security workers secure digital assets, they will also need to create a solid body of research-based evidence to support the practices they recommend. In this project, you read and summarize a research project that addresses Web server security.

Go to *www.sans.org/reading_room/whitepapers/webservers/comparative-study-attacks-corporate-iis-apache-web-servers_33734*. Read the entire article, including the appendices. Answer the following questions:

1. What was the purpose of the research project?
2. What did the research demonstrate about the relative security of the IIS and Apache software?
3. What do you think detracted from the credibility of the report?
4. What were the main findings of the study?
5. What did the results of the study imply about the frequency of automated attacks versus the frequency of manually controlled attacks against Web servers?
6. Based on the results of the study, what policies and procedures would you recommend for Web server security?
7. Based on this study, which operating system is easier to attack: Windows or Linux?

Case Project 12-2: A Recently Discovered Web-based Attack

Use available resources to find a Web-based attack that has been discovered in the last two years. Prepare a one- to two-page report that summarizes the attack and includes the following information:

- The name of the attack
- Its technical mechanism and features
- The presumed goal of the attack
- Severity in terms of risk and damage
- Likely targets
- Known security controls

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 13

Security Policy Design and Implementation

After reading this chapter and completing the exercises, you will be able to:

- Describe the system development life cycle
- Explain the fundamental concepts of risk analysis
- Describe different approaches to risk analysis
- Explain the process of risk analysis
- Describe techniques to minimize risk
- Explain important concepts in security policies
- Identify categories of a security policy
- Define incident handling procedures

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Before undertaking any project, whether you are painting a room in your house or creating a company, it makes sense to imagine the result that you want. The more clarity and detail that you have for your vision of the finished product, the more efficiently and economically you can complete the project. In an organization, tasks are carried out by employees who perform business processes. Typically, these employees work in departments that are responsible for specific functions, such as sales, marketing, research, and production. At a high level, management is responsible for ensuring that the departments are meeting the organization's goals.

To manage business processes effectively, departments and organizations must define their goals clearly. A mission statement is a high-level description of the goals that an organization wants to accomplish. For example, the following excerpt is from the mission statement of Amazon, the Internet sales company:

> *To continue to offer quality products and services using the best technology available and at a reasonable price. This results in highly loyal customers, while maintaining shareholders' interest and company profits in mind. We also want to expand geographically, increasing the number of customers and to keep improving our main competitive advantage—infrastructure.*

Organizations often create value statements to define guidelines for meeting their mission. For example, here is Microsoft's value statement:

> *As a company, and as individuals, we value integrity, honesty, openness, personal excellence, constructive self-criticism, continual self-improvement, and mutual respect. We are committed to our customers and partners and have a passion for technology. We take on big challenges, and pride ourselves on seeing them through. We hold ourselves accountable to our customers, shareholders, partners, and employees by honoring our commitments, providing results, and striving for the highest quality.*

While these statements include specific values and are useful as guides, they do not define how to carry out operations or how to measure the degree to which a company is meeting the terms of its mission and value statements.

Policies are detailed documents that outline conditions, tasks, and measurements within an organization; policies also define conditions and actions that are not permitted. An example of a simple policy would be the requirement that all users who access a database must be authenticated. Other specific details for carrying out the policy would be defined by a procedure document. For example, the authentication policy would reference procedure documents that define acceptable authentication protocols, unacceptable protocols, and the type of authentication used, such as usernames and passwords, biometrics, or smart cards.

As you might imagine, large, efficient organizations have thousands of documents that define policies and procedures. These documents are not created solely to validate mission statements. Many organizations have legal and regulatory standards that mandate how various processes must be undertaken and accomplished. The organization must meet these standards and document compliance. In many organizations, meeting these standards often requires the involvement of the information technology (IT) department, particularly in the area of information security.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Understanding the Security Policy Life Cycle

The development of a security policy follows a life cycle similar to that of software development and many other important long-term projects. A parallel example is a country's foreign policy. A nation's government might have a thorough approach for dealing with political, social, and economic situations in foreign countries, but any number of sudden changes can challenge that approach very quickly. The election of a new prime minister, the passage of a new law, a labor strike, the looming outbreak of hostilities, or the capture of a spy can trigger the need for an immediate change in foreign policy. Similarly, some software products seem to be in a state of constant flux. Updates, patches, service packs, and new versions are released on a regular basis.

The same is true of a security policy. The ever-changing landscape of information security threats makes it unwise to think that a security policy is ever truly complete. On a daily basis, IT staff attempt to guard their companies against the emergence of new attacks and threats.

In a general sense, the life cycle for developing a security policy, software, or foreign policy follows the system development life cycle shown in Figure 13-1. Later in this chapter, you will learn more about each of these four phases, but first you will learn about the goals of each phase.

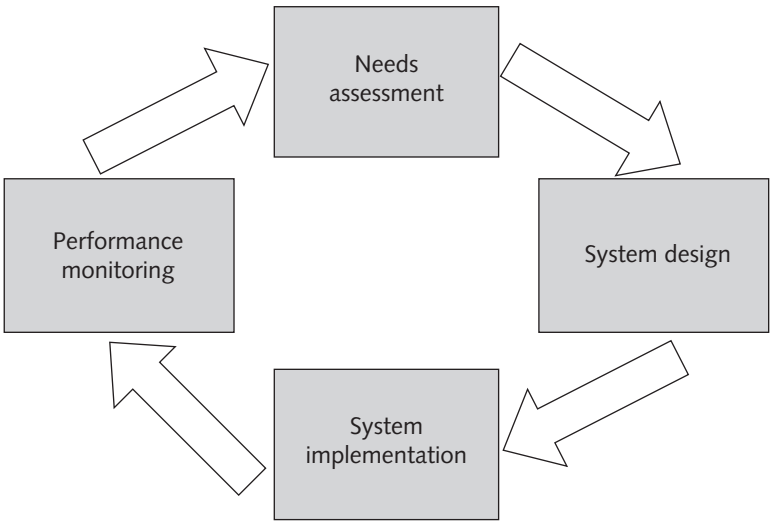

© Cengage Learning 2014

Figure 13-1 The system development life cycle

Needs Assessment

When developing a system or beginning a project, the first important question for an organization is to determine what it needs. The purpose of the system or project must be made clear. Inherent in this question is the issue of how to measure the project's success. Not all system development projects lend themselves to objective assessments, but some standards for success must be established.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

System Design

Once the goal is clear, an organization can begin to create a system that addresses its needs. This phase includes planning. It is critically important to incorporate essential system elements at the beginning of a project rather than having to add them later. For example, it is very difficult to add security components to software after the code has been written; it is much easier to incorporate security as the code is written. Similarly, it is much easier to build legal and regulatory standards compliance into security policies than to rework the policies afterward. Thus, all stakeholders in a project should have input during the planning stage.

Once planning is complete, construction can begin. Whether the project requires writing code or creating security policies, a system of checks and balances should be put into place. For example, team members who are not directly involved in certain portions of a project should check the work of their colleagues in those areas.

System Implementation

Before a system is implemented, user training is often required. Training can be done in stages and using different processes. For example, training for top-level executives might be conducted earlier than training for front-line employees. Similarly, the depth and type of training needed may depend on job function. Typically, the time spent in training and the depth of that training is inversely proportional to an employee's position in the organization. Top-level managers receive less training than mid-level managers; department-level supervisors receive more training than mid-level managers but less than front-line workers.

Depending on the project, a system might be implemented in a pilot phase and activated only with a limited scope, or a system might be rolled out completely and at once. New systems, including security policies, are often rolled out in stages, one department at a time, so that monitoring and needs assessment can proceed without affecting the entire organization. Department by department, the system is implemented.

Performance Monitoring

The monitoring phase can get lost in the press of other duties. Once a system is in place, it is easy to forget about it and move on to the next challenge if the system is not causing any obvious problems. In the case of security policies, this approach can be very dangerous. You may have created a policy for data backups that works well, but if you are not aware that software patches on the company systems have caused conflicts with the automatic backup software, you may be disappointed when you try to restore an important file that has been corrupted. If you are unaware of a change in a regulatory requirement or you missed the news that the accounting department modified a business process and thus subverted a security control inadvertently, your next external audit might be a shocking experience.

It is extremely important to the success of a project to monitor system performance continually. In the case of a security policy, you need to ask several questions: Are any of the assumptions made while developing the policy no longer true? Have new developments required modifications to the policy? Are employees compliant with the policy? Are managers enforcing compliance?

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

After answering the preceding questions and others, you may find that new needs must be addressed in the policy. Thus, you return to the needs assessment phase and the cycle begins again. Systems development, including security policy development, is an iterative and continual process: It cycles through needs assessment, system design, system implementation, and performance monitoring.

Examining the Concepts of Risk Analysis

Before you learn about risk analysis, you should know some of its terminology. For example, an **asset** is a person, thing, or idea that supports the company's mission. Employees, servers, data, and intellectual property are all examples of assets. A **threat** is a person or occurrence that could damage an asset. Examples of threats include hackers, user errors, and acts of nature. A **vulnerability** is a weakness or an exposure that can make an asset more susceptible to risk. An unpatched Web server is a vulnerability because it is a known weakness that can be exploited, but even a patched Web server can create a vulnerability in an organization because an asset is exposed to untrusted systems on the Internet.

Risk is associated with every situation. You take a risk by crossing a street and by opening a business. On the other hand, you might take a risk by *not* crossing the street or creating a business. For example, your business idea might be implemented successfully by someone who acts before you do. **Risk** is the probability that a threat will cause damage to an asset. Because all situations carry risk, there is no situation in which security is perfect.

Your first task when developing a security policy is to assess the risk to your employees, your network, and your databases of customer, job, and personnel information. Your ultimate goal is not to reduce risks to zero, but to devise ways to manage risks in a reasonable fashion. This process, called **risk analysis**, determines the threats an organization faces, the resources that are at risk, and the priority that should be given to each resource. Risk analysis is the first step in formulating a **security policy**, a statement that specifies what defenses should be configured to block unauthorized access to an organization's assets, how the organization will respond to attacks, and how employees should safely handle the organization's resources to prevent loss of data or damage to files.

Risk analysis is an iterative process with a life cycle similar to that shown in Figure 13-1. Figure 13-2 shows the risk analysis life cycle. During the enforcement and monitoring phase, it is common to use insights gained to modify the security policy even as the risk analysis continues.

Many companies, particularly smaller ones, overlook security policies and risk assessments in the process of developing network defenses and countermeasures. When you ask whether a security policy is in place at a smaller company, you may hear responses such as, "*What* security policy?" or "We started working on one three years ago, but it never got anywhere." In the rush to cope with everyday business demands, security policies can easily be overlooked. One of your first tasks as a new security manager might be to convince other managers and employees that they need to develop a risk analysis and security policy cycle for the organization.

The following sections lay the groundwork for understanding risk analysis, which is the first step in developing a network security policy. You learn the fundamental concepts of risk

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

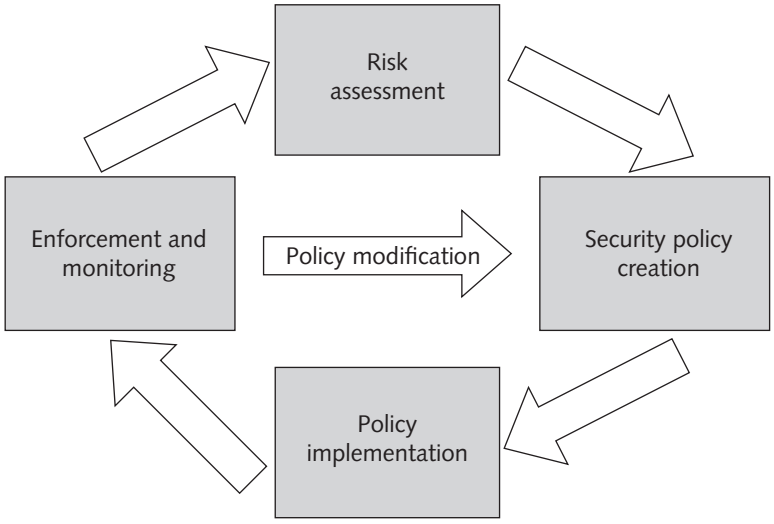

© Cengage Learning 2014

Figure 13-2 The risk analysis life cycle

analysis, different approaches to conducting risk analysis, principles for making risk analysis an ongoing process rather than an isolated occurrence, and ways to analyze the economic effect of threats.

Risk Analysis Factors

In terms of a network connected to the Internet, risk analysis should encompass computer hardware, software, and data warehouses—storehouses of valuable customer, job, and personnel information that a company needs to safeguard.

The following sections describe the six factors needed to create a risk analysis:

- Assets
- Threats
- Probabilities
- Vulnerabilities
- Consequences
- Security controls

Assets Assets in an organization play a central role in risk analysis—after all, they are the hardware, software, and informational resources you need to protect by developing and implementing a comprehensive security policy. You are likely to encounter four types of assets:

- *Physical assets*—Equipment and buildings in the organization
- *Data assets*—Databases, personnel records, customer or client information, and other data the organization stores and transmits electronically

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *Software assets*—Server programs, security programs, and other applications used to communicate and carry out the organization's typical activities
- *Personnel assets*—People who work in the organization as well as customers, business partners, contractors, and freelance employees

Some assets are tangible objects, such as computers. Other assets are intangible, such as a company's reputation and the level of trust it inspires in its customers. You might consider other assets to be essential business concepts, such as confidentiality, integrity of information, and availability of resources. The most valuable information, such as database contents that need to be confidential and accurate, is what you should focus on first. Data might be the most important asset to continuing business operations, but it is also the most difficult asset to evaluate. In fact, listing every single asset you have can be difficult. You might be able to analyze only the most important assets in detail.

Threats Threats are events and conditions that could potentially occur, and their presence increases risk. Some dangers are universal, such as weather-related disasters. Others are more specific to your system, such as a server storing a customer database; an attacker could exploit the server to gain access to the system. Other examples of threats include the following:

- *Power supply*—The power supply in your area might be unreliable, making your company subject to brownouts, blackouts, and sudden surges called voltage spikes.
- *Crime rate*—If you work in a high-crime area or other offices in your area have been burglarized, your risk increases.
- *Facility*—If your building has old wiring that is prone to fluctuations or has insufficient fire suppression, the risk of fire damage increases.
- *Industry*—If your organization operates in a highly competitive industry or one that requires high security, a security breach could result in litigation or major loss of revenue or even force the business to close.

The seriousness of a threat depends on its probability of occurring, as explained in the following section.

Probabilities Geographic or physical location, habitual factors, and other factors affect the **probability** that a threat will occur. A geographic factor might include earthquakes, which are common threats in some regions. Physical location might also influence threat probability; an example would be an electrical problem in the building that houses your systems. Habitual factors could be poor security practices, such as employees keeping written passwords exposed near their computers, that increase the probability of a security breach. These factors are a large part of what risk assessment seeks to uncover. Risk analysis evaluates each factor and rates its potential impact or exposure.

Your **exposure** to risk increases if your organization has one or more factors that increase threat probabilities. For example, if you live in a part of the country with frequent severe storms or floods, the threat of weather-related damage increases. If several disgruntled employees had access to sensitive information and have just been fired, the probability of losing that information increases unless you take steps to protect it before they leave. If your office has an alarm system wired to a security service, the probability of burglary is

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

reduced. The probability of threats occurring is often assessed and recorded in general terms, as shown in Table 13-1.

Threat	Probability
Earthquake	Medium
Fire	Low
Flood	High
Attack from the Internet	Very high
Virus infection	Very high
Employees giving out information	Low

© Cengage Learning 2014

Table 13-1 Sample threat probabilities

Vulnerabilities Vulnerabilities are situations or conditions that increase the probability of a threat, which in turn increases risk. Examples include connecting computers to the Internet, keeping computers in open areas where anyone can use them, and installing Web servers outside the corporate network in the vulnerable demilitarized zone (DMZ).

You can easily think of examples of vulnerable situations that affect networked hardware and software. Some common flaws involve operating system software, but in recent years, application software has also become a prime target of attackers. Even the freeware Web server Apache has fallen victim to security compromises as a result of software flaws. Although some systems have more security flaws than others, remember that every system has flaws. Opening a network to remote users whose computers are unprotected by anti-virus or firewall software can expose the network to intrusions and virus infections. Poorly configured firewalls or packet filters, unprotected passwords, log files that are not reviewed closely or regularly, new intrusion threats that affect wireless networks, and the complexity of modern computer networks all contribute to vulnerabilities that can give attackers an opening.

Consequences Substantial adverse consequences can result from a virus that forces you to take the corporate Web site offline for a week or a fire that destroys computer equipment. You can extend the earlier identification of threats to include ratings that evaluate consequences of those threats, as shown in Table 13-2.

In Table 13-2, the probability of threats has been extended to include a rating of their impact. Ranking these items can be difficult because the severity often depends on a specific virus or your physical location, for example. A flood can pose a serious threat to computers on the ground floor of an office, but not to computers on the tenth floor.

Besides the consequences of getting a system back online after an attack, cost impacts and other effects can be more difficult to anticipate, including insurance claims, police reports, shipping or delivery charges, and the time and effort needed to obtain and reinstall software

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Threat	Probability	Consequences
Earthquake	Medium	Significant
Fire	Low	Significant
Flood	High	Minor
Attack from the Internet	Very high	Serious
Virus infection	Very high	Serious
Employees giving out information	Low	Significant

© Cengage Learning 2014

Table 13-2 Probabilities and consequences of threats

or hardware. The actual cost of an incident is usually much higher than the cost of replacing equipment and restoring data. When you go to management to justify investments in security, estimating the cost of the investment and its benefit to the company is vital. This estimate is commonly called a **cost-benefit analysis**. The most critical numbers you want management to understand are the actual costs paid per year by the company because of security incidents. The benefit is the amount per year saved by preventing these incidents.

Security Controls Security controls are countermeasures you can take to reduce threats, such as installing firewalls and IDPSs, locking doors, and using passwords and encryption. These measures interact with each other to help manage risk. When deciding how to manage risk, you must first identify and classify the risks. Next, you determine priorities of threatened assets, and then you determine whether to accept, transfer, or mitigate the risk.

An asset has an associated amount of risk. Threats and vulnerabilities increase the risk; countermeasures work to reduce risk. **Residual risk** is the amount left over after countermeasures are implemented; a risk never actually equals zero. Figure 13-3 illustrates this process.

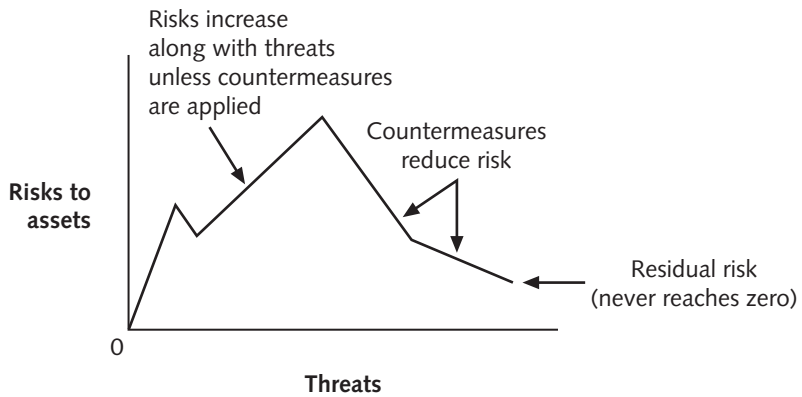

© Cengage Learning 2014

Figure 13-3 Countermeasures reduce but never eliminate risk

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Risk Analysis Methods

After you have studied all the risk analysis factors that affect your organization, you have the building blocks needed to prepare a risk analysis. You can use different methods of risk analysis to create a security policy and then evaluate how well the policy is performing so that you can update and improve it. The following sections describe the two methods you are most likely to use.

Survivable Network Analysis

Survivable Network Analysis (SNA) is a security process developed by the CERT Coordination Center. SNA starts with the assumption that a computer system or network will be attacked. It leads you through a four-step process designed to ensure the survivability of a network if an attack occurs. **Survivability** is the ability to continue functioning during attacks, system faults, accidents, or disasters.

Survivability focuses on a network's essential services, assets, and critical capabilities and depends on four key properties of a network:

- *Resistance*—The capability of a system to repel attacks
- *Recognition*—The capability to detect attacks when they occur and to evaluate the extent of damage and compromise
- *Recovery*—The capability to maintain essential services during an attack and restore all services afterward
- *Adaptation and evolution*—The capability to improve system survivability based on knowledge gained from attacks

The study of a network's survivability builds on other concepts related to risk analysis, including **fault tolerance**, which is the ability of an object or a system to continue operations despite a failure, such as a system shutdown. Safety procedures, security systems, and ongoing testing are also important to survivability. Most software products are not designed with survivability in mind, which is why survivability studies can be valuable. Instead, software is often designed to work for a certain number of users or for a certain amount of information until it is replaced by new and improved software.

The steps in SNA are as follows:

- *System definition*—First, you create an overview of the system's organizational requirements. You analyze system architecture while accounting for its hardware components, software installations, databases, servers, and other computers that store information.
- *Essential capability definition*—You identify a system's essential services and assets that are critical to fulfilling your organization's missions and goals.
- *Compromisable capability definition*—You design situations in which system intrusions occur and then trace the intrusion through your system architecture to identify what can be accessed and what sorts of damage can occur.
- *Survivability analysis*—You identify potential points of fault in the system—integral components that can be compromised. You then make recommendations for correcting the points of fault and for improving the system's resistance to intrusions and ability to recover from attacks, accidents, and other disasters.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The emphasis is on creating an ongoing process rather than attempting to create a final, static security posture. You might start with better password management, then upgrade the system to encrypt critical data, and then install software that filters out potentially harmful e-mail so that the system's ability to survive improves continually.

Threat and Risk Assessment

Threat and Risk Assessment (**TRA**) approaches risk analysis from the standpoint of threats and risks to an organization's assets and the consequences if those threats and risks occur. Like SNA, TRA has four steps:

- *Asset definition*—You identify the software, hardware, and information you need to defend.
- *Threat assessment*—You identify the kinds of threats that place the asset at risk, including vandalism, fire, natural disasters, and attacks from the Internet. Threat assessment also includes an evaluation of the probability and consequences of each threat.
- *Risk assessment*—You evaluate each asset for any existing safeguards, the severity of threats and risks to each asset, and the consequences of the threat or risk taking place. The combination of these factors creates an assessment of the actual risk to each asset.
- *Recommendations*—Based on the risks and current safeguards, you make recommendations to reduce the risks. These recommendations should then become part of a security policy.

Table 13-3 lists ratings you can assign to describe the probability of threats occurring.

Rating	What it means
Negligible	Unlikely to occur
Very Low	Likely to occur only two or three times every five years
Low	Likely to occur within a year or less
Medium	Likely to occur every six months or less
High	Likely to occur after a month or less
Very High	Likely to occur multiple times per month or less
Extreme	Likely to occur multiple times each day

© Cengage Learning 2014

Table 13-3 Describing the probability of threats

After rating the severity of a threat or risk, you evaluate the consequences if it actually occurs. Table 13-4 shows one method of describing consequences.

After evaluating the level of threats to assets and describing the consequences of the threats occurring, you can combine these two ratings to develop a risk analysis of each asset, as described in the following section.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Description	Consequences
Catastrophic	Threatens the continuation of the program and causes major problems for customers
Major	Threatens the continuation of basic program functions and requires intervention by senior-level management
Moderate	Does not threaten the program; however, the program could be subject to major review and modification of operating procedures
Minor	Could threaten the program's efficiency or effectiveness but can be handled internally
Insignificant	Can be handled by normal operations

© Cengage Learning 2014

Table 13-4 Describing consequences

The Risk Analysis Process

Risk analysis is not a one-time activity used to create a security policy. Rather, risk analysis evolves to account for an organization's changing size and activities, the progression to larger and more complex computer systems, and new threats from inside and outside the corporate network.

The initial risk analysis is used to formulate a security policy; the policy is then enforced and security is monitored. New threats and intrusion attempts create the need to reassess the risk an organization faces.

General Activities to Follow Risk analysis is a group of related activities that typically follow this sequence:

- *Holding initial team sessions*—First, hold meetings to get groups of workers together in one place. Conduct interviews or hand out questionnaires to collect pertinent information. It is especially important to talk to all managers to set the objectives and scope for the risk analysis, schedule how long the project should take, and identify the important people you need to interview.
- *Conducting asset valuation*—After you determine the scope of the risk analysis, you need to identify assets to protect and determine their value. This activity can be classified as subjective or speculative. If the activity is subjective, you are assessing the impact of losing assets that might not be tangible, and you should use your best judgment or solicit opinions from other qualified employees. If the activity is speculative, you are trying to determine whether information might fall into the hands of unauthorized people and estimating the cost of recovering the information. Personal interviews with managers can help you determine a realistic assessment.
- *Evaluating vulnerability*—You investigate the levels of threat and vulnerability in relation to the value of the organization's assets. Ask IT staff to evaluate the threat of virus attacks or other intrusions on a scale of one to five, for instance.
- *Calculating risk*—After you have determined asset values and the vulnerabilities that threaten those assets, you can calculate risk. Usually, a numeric value is assigned.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

For instance, 1 might represent a low-level baseline security need and 7 might represent a high-security priority.

Remember that the first step in risk management and designing and implementing a security policy is to gain the backing of upper management, such as the board of directors, the CEO, and the CIO. This backing makes it easier to get cooperation from colleagues, other departments, and other employees. Convincing managers that the organization needs a security policy is not difficult if you focus on how security affects profitability. You must be able to support the data you provide, and you must not exaggerate the potential costs of a security incident.

Analyzing Economic Impacts

An important part of conducting a risk analysis is preparing estimates for the financial impact of losses. If you are familiar with statistics, you can use several different models for estimating impacts. You can also use a software program to help you prepare reports that substantiate your estimates and produce charts and graphs to support your figures. Project Risk Analysis by Katmar Software, for example, gives you a structure for listing hardware and software assets in your organization (see Figure 13-4). You will use this program in Hands-On Project 13-1.

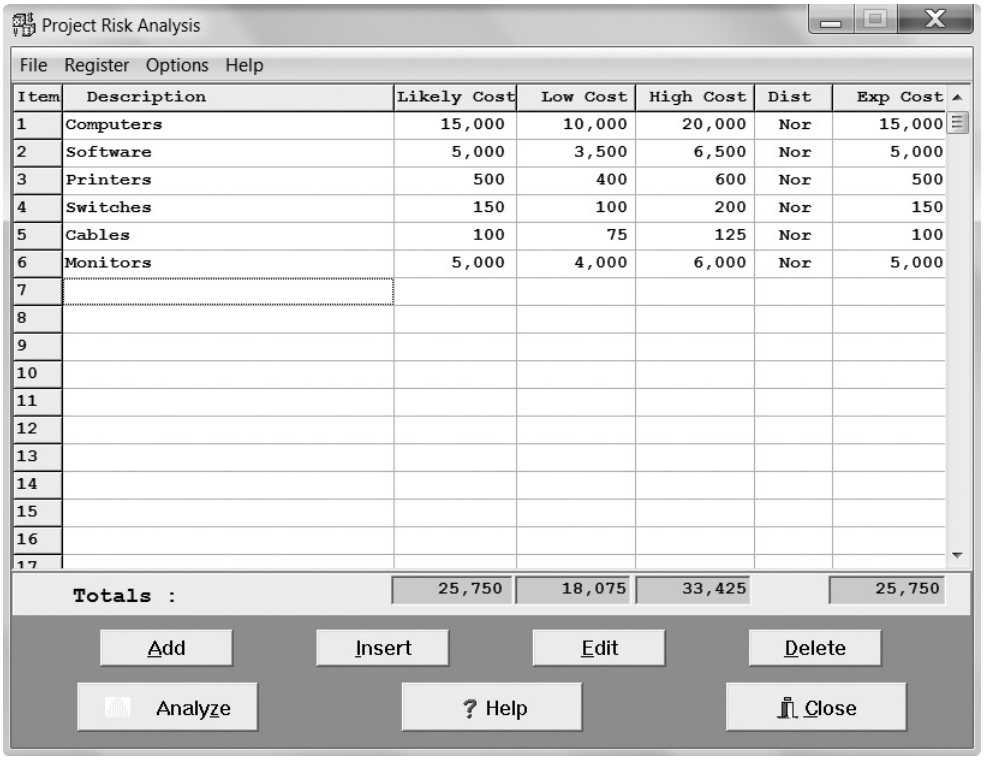

Source: Katmar Software

Figure 13-4 Project Risk Analysis offers a structure for making cost estimates

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

With this program, you can make cost estimates by using a variety of statistical models. If you are unfamiliar with statistics, you can try the simplest model, which uses the following estimates:

- *Likely cost*—The most realistic estimate of the money needed to replace the item
- *Low cost*—The lowest dollar amount for replacing the item
- *High cost*—The highest dollar amount for replacing the item

When you create a record of an asset in Project Risk Analysis and estimate its replacement cost, you enter the preceding values using the Normal distribution setting, as shown in Figure 13-5.

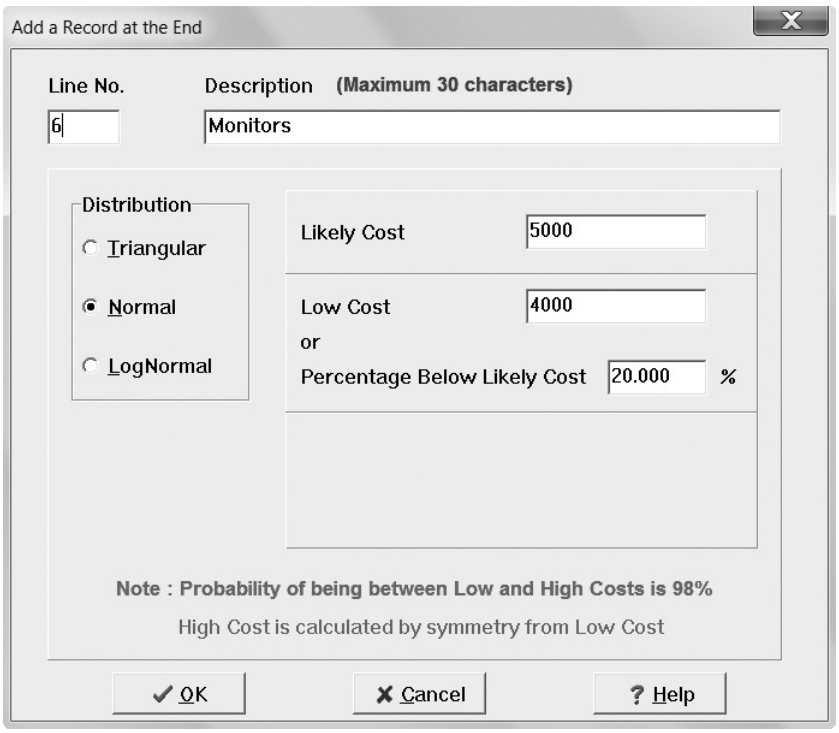

Source: Katmar Software

Figure 13-5 Entering values for replacement costs

One advantage of using a risk analysis program is that you can analyze cost estimates and present them in a report format (see Figure 13-6). In addition, these programs can quickly calculate the mean cost of replacing hardware, software, or other items.

Project Risk Analysis performs calculations by using a statistical formula called a **Monte Carlo simulation**—an analytical method that simulates a real-life system by randomly generating values for variables. The charts and reports these programs create are valuable documentation for preparing a risk analysis; the visual elements can have significant impact when presented to managers.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

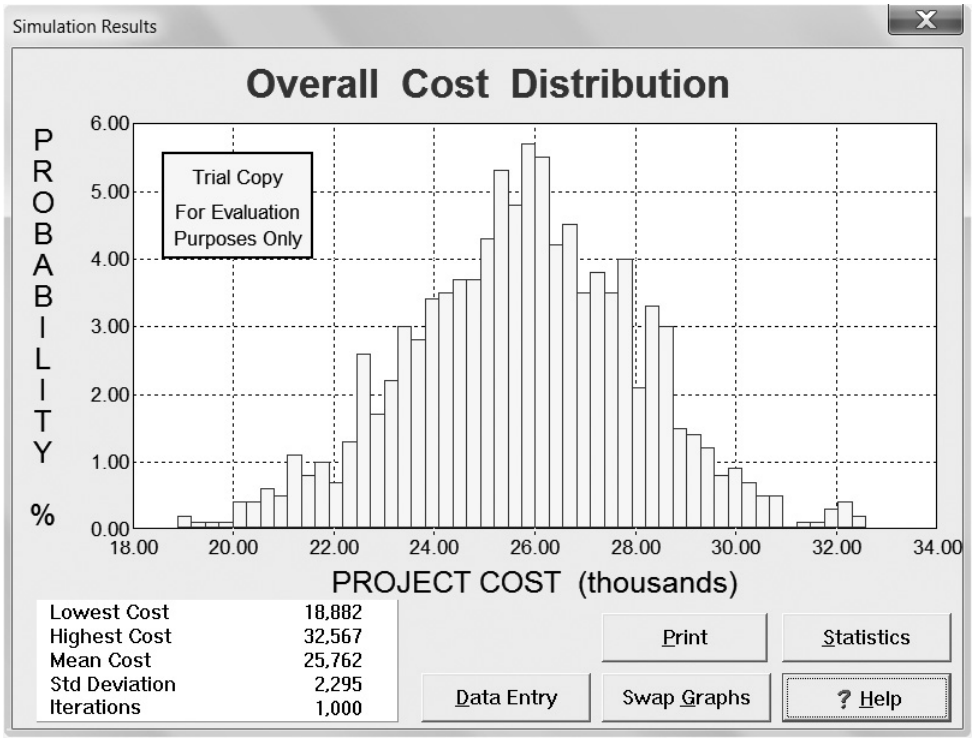

Source: Katmar Software

Figure 13-6 Graphical reports in risk analysis software

Techniques for Minimizing Risk

After analyzing the level of risk to hardware and software assets in your network, you can recommend safeguards for minimizing the risk. **Risk management** is the process of identifying, choosing, and setting up countermeasures for the risks you identify. The countermeasures you describe are statements you should incorporate into your security policy. In the following sections, you learn about important points for deciding how to secure hardware, how to secure information databases in your network, how to conduct routine analysis, and how to respond to security incidents.

Securing Hardware Your company's physical computing assets—the hardware devices that keep data flowing throughout the network—are the most obvious objects you need to identify. You must decide how to protect your hardware. First, think about obvious types of physical protection, such as environmental controls to keep machines safe in hot temperatures and fire protection systems. Then consider whether you need to lock up all hardware in your organization or use theft protection only for servers. Place your servers in a locked room with an alarm system so that unauthorized employees cannot access them; this equipment should never be left in the open.

Be sure to pay special attention to laptop computers and handheld devices in your organization. These machines can be lost or stolen easily, and any proprietary information on them could be compromised. Such incidents happen regularly and often with serious consequences. To alleviate

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

this problem, be sure to install startup passwords as well as screen-saver passwords on all mobile devices. Experienced thieves can circumvent these passwords, of course, but at least they make it more difficult to access files. In addition, you can use encryption to protect data.

Conducting a Hardware Inventory Make a list of servers, routers, cables, computers, printers, and other hardware the company owns. Be sure to include your company's **network assets**—the hardware and software that enable employees to communicate with one another and other computers on the Internet. Create a topology map that shows how the devices are connected and includes an IP allocation register, such as the one in Figure 13-7.

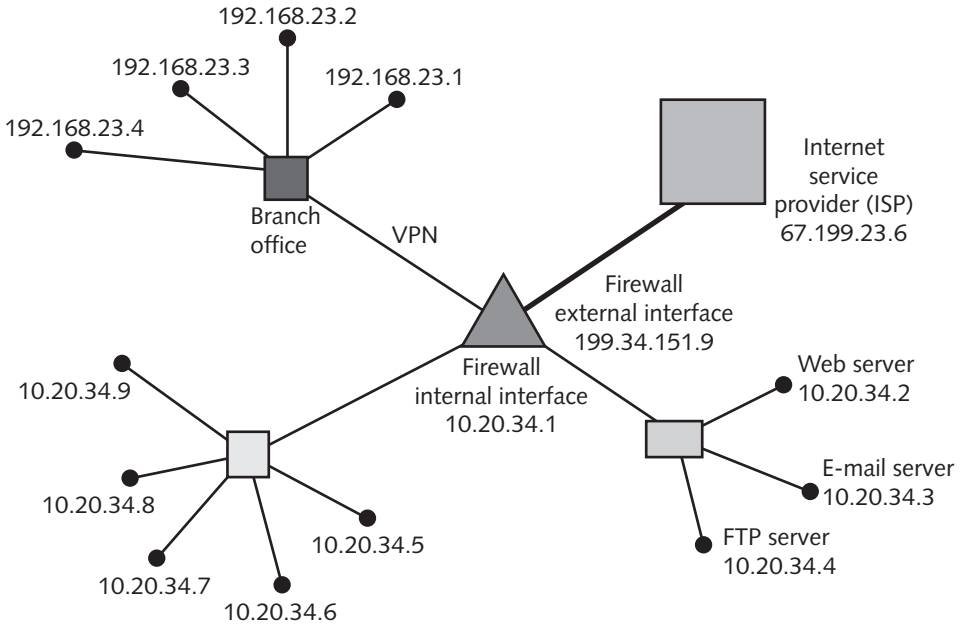

© Cengage Learning 2014

Figure 13-7 A topology map can supplement a hardware inventory

Ranking Resources to Be Protected When listing physical, electronic, network, and system assets, assigning a value to each object is helpful. The value can be an arbitrary number; the object is to rank resources in order of importance so you can focus your security efforts on the most critical resources first. The team that helps you prepare your security policy will probably determine that data is more important than the devices that store the data.

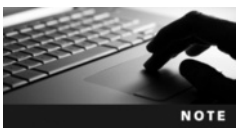

Although your rankings might seem arbitrary, deriving them with the cooperation of your organization's top management is helpful. Developing a lengthy list of resources and rankings on your own without input from managers is likely to result in extensive revisions. You will get better results if you submit a list of resources to managers and ask them to develop their own rankings. Also, ask them to consider the cost of replacing software and computers you have listed. Suggest that they rank assets on a scale of 1 to 10.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Using Encryption One of the best and least expensive methods of securing data is to apply encryption that locks it down. Encryption does not prevent intruders from accessing or viewing encrypted data or even prevent its theft. However, encryption can prevent data from being exploited. The following sections describe several areas in which using encryption could be helpful in minimizing the risk of sensitive data being compromised.

- *Mobile computers*—Thieves and attackers are targeting laptops, notebooks, and smartphones more often as sources of sensitive data, so they should be considered high-security devices. Mobile computer theft can result in large-scale unauthorized access to sensitive corporate and personal data. This damage can be prevented easily by using physical security, password protection, and data encryption on mobile devices.
- *Removable media*—When you are considering data vulnerability, remember that storage disks, tapes, CDs, flash drives, and memory chips are much like mobile computers. Removable media are easy to transport and conceal, and they can contain large quantities of data. As with mobile computers, using physical security, password protection, and data encryption can protect sensitive data from being compromised.
- *Data transfers*—Moving data over wired and wireless media is common. Devices that share data often use encrypted authentication methods to establish a communication link. However, after communication paths have been established, data might be transferred in cleartext, making it vulnerable to theft and exploitation.

Securing Information After you have decided on safeguards for your hardware, you need to determine how to protect your company's **electronic assets**—word-processing documents, spreadsheets, Web pages, and other documents on your network computers. Logical assets include e-mail, any records of instant messaging conversations, and log files compiled by firewalls and IDPSs. Data assets include personnel, customer, and financial information that your company needs to protect.

Many companies now conduct all or part of their business operations on the Internet. If your organization conducts e-commerce, you need to strike a balance between making transactions easy for customers and keeping information confidential.

One way to protect information customers send via the Internet is to isolate the information so that attackers cannot access it. In your security policy, you might want to state that to minimize the risk of attackers stealing credit card information and other critical customer data, your company needs to move information to a computer that is physically isolated from the Internet. You can configure backup software to save critical files in isolated locations automatically on a nightly or weekly basis. You can also use the following measures to protect information:

- *Encryption*—By encrypting data, you can protect it as it passes from one network to another. Encrypted data cannot be read if it is intercepted or captured.
- *Message filtering*—This measure keeps potentially harmful messages from entering the network from the outside.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *Data encapsulation*—The data in packets can be encrypted so that the packets are encapsulated (or "wrapped") for extra protection.
- *Redundancy*—By providing redundancy through backup systems, you ensure that databases and other stores of information remain accessible if primary systems go offline.
- *Backups*—Making systematic and periodic backups of information on the network is one of the most basic and important ways to protect that information.

Do employees at your organization handle confidential, proprietary, or private information? If so, this information needs to be covered in the security policy. Safeguards are needed to inform employees of their responsibilities for protecting sensitive information. To minimize risks, you could specify the following measures in a security policy:

- Never leave company-owned laptops or handheld devices unattended.
- Always protect information on corporate devices with passwords.
- Encrypt any confidential information.
- Use passwords to protect all job records and customer information.
- Restrict access to personnel information to human resources staff and upper management.

You need to make sure that all employees read and understand the policy. You might consider distributing the policy in the form of a manual issued to all new employees. You can publish the policy on the company Web site for current employees to review. You can also require employees to sign a statement that they have read the policy and understand their responsibilities.

Conducting Ongoing Risk Analysis Risk analysis is an ongoing operation. A company changes constantly in terms of the information it handles, its number of customers and employees, and the number of computers on its network. Risk analysis should be conducted on a regular basis despite common obstacles, such as the indifference of IT staff and employees, heavy workload in critical areas, and lack of available personnel to do the evaluation.

Deciding how your organization should perform routine risk analysis starts with the following questions:

- How often will risk analysis be performed? A common approach is to schedule the analysis every year at budget time. However, you can keep up with new threats more effectively by conducting a risk analysis every six months.
- Who will conduct the risk analysis? The same professionals who manage security for the organization are the ones who should participate, along with accounting or bookkeeping staff.
- Do all hardware and software resources need to be reviewed every time? You might not need to conduct a new risk analysis for every asset you have; you might decide that you need to examine only assets that have increased or changed substantially since the last review.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

The calculations and evaluations in risk analysis require subjective evaluations of how much an asset is worth and how valuable it is. Human emotions can influence evaluations, so many companies do not allow employees to perform these calculations manually. Because calculations are often complex, using risk analysis software can make the process easier as well as more objective.

Examining the Concepts of Security Policies

You may sometimes hear people question whether their organization really needs a security policy. You should remind your skeptical coworkers that a security policy is indeed necessary, particularly if the organization falls into one of the following categories:

- Employees work with confidential or proprietary information.
- Damage, theft, or corruption of systems or data could result in severe financial losses that endanger business continuity.
- The organization has trade secrets that are important to its goods or services.
- Employees regularly access the Internet and use e-mail or other means of electronic communication that could be attacked or infected.
- The company is part of an industry that is subject to state and federal regulations for information security and privacy.
- The company uses Internet connections with partner businesses or application service providers (ASPs)—companies that provide Web-based services for a fee.

To help your case for establishing a security policy, you can also provide examples of problems that can occur if a clearly defined policy is not in place:

- A copy editor spent a lot of time at his job surfing the Internet for personal reasons and work-related research. Management could not discipline this employee because it had no policy for what constituted excessive personal use. The employee's supervisor later discovered that the copy editor had downloaded pornography; for this reason, the employee was fired. However, the employee subsequently appealed the dismissal with the Civil Service Board, claiming that he had never been told he could not download pornography. After a hearing, the board ordered him to be reinstated with back pay.
- A clerk who was laid off because of downsizing at an insurance company was hired by a competing company. The insurance company discovered that many of its customer files had been accessed and copied, and that a number of its clients had switched to the competitor after being offered lower insurance rates. The company did not have a policy for protecting its passwords or switching passwords after employees left. The laid-off clerk was able to access her former employer's network and steal files so that her new employer could market itself aggressively to potential customers.

The benefits of a security policy are wide ranging. In general, however, a security policy provides a foundation for an organization's overall security stance. A security policy gives employees guidelines for handling sensitive information, provides the IT staff with instructions for configuring defensive systems, and reduces the risk of legal liability for the company and

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

its employees. Employees who are fired or laid off often sue their previous employers or file a grievance; in these cases, a well-defined security policy can make the difference between the company having to pay damages and not being liable.

To protect security, it is important to formulate a clear policy that states the rights of employees and how they should handle company resources responsibly. All employees should sign the policy when they are hired. In some instances, an employee who works with sensitive or proprietary information should be required to sign a nondisclosure or confidentiality agreement.

A good security policy is comprehensive and flexible; often, it is a group of documents, each with a specific emphasis. The next section discusses general best practices for security policies. Then you examine major considerations that can help you develop an effective set of rules and procedures:

- The decision to purchase cyber risk insurance coverage
- The need to base a policy on a thorough risk assessment
- The need to teach employees about acceptable use of network resources
- The need to specify an employee's privacy rights while on company property and using company equipment
- The need to enable management to set priorities
- The need to help administrators do their jobs
- The need to see a security policy as a vehicle for good risk analyses

General Best Practices for a Security Policy

A security policy for a large corporation might rival the size of a metropolitan phone book and encompass every possible aspect of the business, or it can be a simple document that describes a few fundamental rules for a small company. Regardless of your organization's size, environment, and needs, you need to understand a few basic concepts about building an effective security policy:

- If it is too complex, no one will follow it. In fact, users might circumvent it.
- If it hurts productivity, it will fail.
- It should state clearly what employees can and cannot do on company property and with company equipment. Avoid jargon or complex descriptions, but be as thorough as possible.
- Include general clauses to summarize statements, such as "Employees are not permitted to download games, screen savers, wallpaper, images, video clips, art, *or any other form of multimedia applications or files*." The corporate attorney might need to fine-tune the wording, but the italicized clause covers anything not specifically mentioned. A similar phrase is "*including but not necessarily limited to*."
- People need to know why a policy is important. They are more likely to accept it as necessary if they understand it.
- When developing the policy, involve representatives of all departments, including rank-and-file employees. The benefits are twofold: First, you will design a more

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

accurate and appropriate policy if you tailor it to fit the needs of people who actually use the systems. Second, by involving every level of the company, you give employees a personal stake in the process. This sense of ownership leads to a more involved attitude and better morale, which equates to a more effective and enforceable security policy.

- The policy should contain a clause that describes the consequences an employee could face for violating the policy.
- The policy must have support from the highest level of the company, and that support must flow down through the ranks. If management does not endorse or obey the policy, why should employees?
- Have all employees sign a document acknowledging their understanding of the policy and agreement to abide by it. Some might think this practice conveys a lack of trust, but it is a sensible precaution that serves as an audit trail.
- Keep your security policy updated with current technologies. For example, include guidelines for mobile devices and regulate the use of portable storage devices. The point is not to cover specific devices, but to cover categories of technology. For example, instead of describing policies for cell phones or USB drives, specify policies for wireless devices or portable storage devices. Remove outdated material that no longer applies or that has been integrated into another area. For example, if the policy is updated to include a special section on laptops and portable computing devices, you should remove statements about these devices in other areas of the policy.
- Make certain that your policy directives are consistent with applicable laws. Retaining legal counsel to review your policy draft might be prudent to make sure that all bases are covered and that the policy does not violate civil rights or other laws. For example, in the United States, it is illegal to forbid employees from publicly protesting a company's actions unless they agree to that provision. Even then, this clause could be challenged based on the First Amendment of the U.S. Constitution. Also, make sure that the security policy does not conflict with other corporate policies.

Although this task might seem overwhelming at first, keep in mind that templates and other free resources are available. You can customize these resources to fit your needs.

Developing Security Policies from Risk Assessment

To start developing a security policy, you identify the assets that need to be protected. You then define threats the network faces, the probability that those threats will occur, and the consequences posed by each threat. Finally, you propose safeguards and define how to respond to incidents.

The information you gather during risk analysis belongs in the security policy, along with a statement of the policy's overall goals and the importance of employees reading the policy and following its guidelines. The penalties for violating the policy should be stated near the beginning of the document. In litigation, the company attorney then can argue that the employee could not have missed the penalty section and had to know that violating the policy could result in disciplinary action. Most policies list termination as a possible penalty, mainly to protect the company from wrongful termination lawsuits or similar grievances. As you can see, an important function of a security policy is to reduce legal liability.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

After the security policy is implemented, its effectiveness must be monitored. Some ideas that sounded good in theory might not work well in practice. Like any large system with interdependencies, an organization-wide security policy requires periodic revision. Sometimes an immediate change is required in response to a problem, but other revisions can be made on a regular cycle.

Remember that a security policy is a fluid document that must adapt to new circumstances to remain effective. As problems, intrusions, or incidents occur, you can update the policy as needed; the sequence of policy design, implementation, monitoring, and reassessment is an ongoing cycle.

Teaching Employees About Acceptable Use

The issue of trust is an integral part of a security policy. The policy needs to define which employees to trust and the levels of trust that all employees should receive.

In reality, organizations must achieve a balance between trusting employees and issuing orders. Employees must be able to use their computers to communicate and be productive. By placing too little trust in people and regulating everything they do in an excessively rigid fashion, you hamper their work, hurt morale, and increase the odds that employees will circumvent security safeguards.

Outlining Penalties for Violations

The cornerstone of many security policies is the acceptable use policy, which defines how employees should use the organization's resources, including the Internet, e-mail, and software programs. More important, the policy should spell out what constitutes unacceptable use, such as downloading or viewing objectionable or offensive content, using company equipment for personal business, and removing company property (including any digitized information) without specific permission.

Policy declarations exist to guide management and employees in the proper use of corporate resources. Penalty clauses exist so that companies can discipline employees whose computer activities interfere with productivity. The policy should also contain penalty guidelines so that employees cannot claim ignorance. For example, the first offense could lead to a warning, the second offense could require a notice in an employee file or a probation period, and the third offense could result in terminating the employee. In extenuating circumstances, an employee could be given another chance. The idea is to establish flexible methods of punishment that can be applied at management's discretion but that protect the company from litigation. Wrongful termination lawsuits are common and can result in high costs for a company. Penalty clauses, if worded correctly, can protect companies from such lawsuits. Again, legal counsel should review the security policy.

Criminal Computer Offenses

If the security policy violation is a criminal offense, such as possessing child pornography, law enforcement officials must be notified and the investigation is turned over to them.

In this situation, you need to be aware of some pitfalls. After an investigation is turned over to the police, the U.S. Constitution's **Fourth Amendment** protections for search and seizure apply. Your responsibilities as administrator change when public entities such as police and

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

government become involved and company information becomes evidence. You face a tremendous burden of liability if you assist the police without being served with a **subpoena** or **search warrant**. A subpoena is a court order demanding that a person appear in court or produce some form of evidence. A subpoena specifies what evidence is required, such as personal testimony, paper documents, or data files from the suspect's computer. A search warrant is similar to a subpoena, in that it is issued by a court and compels you to cooperate with law enforcement officers. It also describes the location to be searched and specifies the evidence that officers are allowed to take.

Search warrants and subpoenas were developed in response to the Fourth Amendment, which protects U.S. residents against illegal search and seizure. You can be sued for violating a person's constitutional rights if you provide evidence without being served with a warrant or subpoena. Providing assistance without being served is known as becoming a "de facto agent of law enforcement." You are acting under their direction and aiding them in an investigation, so you are bound by the same procedures and laws that bind investigators. When turning over a case to police officers, you can provide all the evidence you have gathered to that point, but you cannot continue to investigate from that point unless you are ordered to do so. If you have been served with a subpoena or officers arrive to carry out a search warrant, you must comply to remove any potential liability.

Also, employees have an expected right to privacy. The lines between Fourth Amendment protections and privacy expectations versus company-owned property can become blurry unless the security policy spells out an employee's expectations for privacy while using company equipment at work. Your security policy must state clearly that company equipment and all digital information on it or accessed through it are the exclusive property of the company, and that the company reserves the right to search and inspect its property at any time. The policy must also state that an employee has no expectation of privacy while using company resources. Unless this expectation is stated specifically in the policy, the company and any parties to an investigation could be liable if evidence is revealed that incriminates an employee. For example, if an employee is terminated for statements made in an e-mail, he could argue successfully that he did not realize the company would read what he considered private e-mail, and that he should not be held accountable for what he said. Some companies might want to have employees sign a separate acknowledgement of this clause.

Enabling Management to Set Priorities

Security policies give employees guidelines they can follow during everyday work activities, but they are also helpful to management. They protect management in disputes with employees or if employees complain about security measures the organization uses.

Security policies do more than assist with dispute resolution, however. They give management a way to identify security priorities for the company. A security policy lists the network resources that managers find most valuable in the organization and that are most in need of protection. It spells out whether the organization allows employees to access the Internet and internal corporate network. In addition, it describes measures the company takes if resources are misused or information is compromised by people outside or inside the company.

When you prepare a project proposal for developing full-fledged security policies or purchasing expensive security-related hardware and software, management will probably ask you to

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

estimate the return on investment (ROI) the company will realize. ROI indicates how long it takes for an investment to pay for itself. Advise management to think about the issue as it relates to ROI but in terms of the cost of doing nothing. Large organizations tend to have employees who connect remotely, and thus are far more vulnerable to security breaches than other businesses. These organizations need to consider three factors:

- The value of information systems and the data in them
- Threats that the organization has encountered and will encounter
- The chances that security threats will result in lost time and money

Security threats can originate from many sources and have a wide variety of signatures, and no single solution can block them all. A comprehensive security plan should not be evaluated strictly in terms of ROI, but as a vehicle for protecting systems and customer information. Remind management that the company could lose significant revenue if attackers steal customers' credit card numbers and post them on the Internet. Such an incident could destroy the company's credibility and expose it to possible criminal and civil litigation. In addition, ask management to estimate the cost to the business if information is unavailable for weeks or months. You should present these scenarios diplomatically and have statistics to back up your points.

Dealing with the Approval Process

Developing a security policy can take several weeks, months, or even longer, depending on the organization. Do not rush the process; take the time to do it right and cover all the bases. A poorly written security policy might be rejected by management.

As part of development, upper management, executives, and other stakeholders need to review and approve the security policy. This process can take several weeks or months. Do not be dismayed if the process takes longer than you expect. For a security policy to work, employees must accept it. The policy might encounter resistance, which is natural for any policy that affects an entire organization. To help gain acceptance for the policy, you can establish a security user awareness program, in which employees are instructed formally about the organization's security strategy.

Feeding Security Information to the Security Policy Team

Any changes to an organization's security configuration should be conveyed to the security policy team. This team can suggest changes to the policy and determine whether new security tools need to be purchased or new security measures need to be taken.

Management's participation and backing can help in amending the security policy. Encourage managers to inform employees that protecting company assets is everyone's responsibility. Provide training to users to make them aware of security issues, and explain why data should be collected and managed securely. Educate users to do their jobs in a secure manner. Listen to employees' concerns. Develop sensible security solutions that allow daily business to be conducted yet provide an acceptable level of protection against risks.

Helping Network Administrators Do Their Jobs

Network administrators can get considerable assistance from security policies if they are responsible for showing employees how to access shared resources, change passwords, sort

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

through e-mail, and perform other functions. A security policy can spell out simple but important information that an administrator would otherwise have to convey personally. For example, the policy might contain the following rules:

- Users are not allowed to share accounts with other employees, visitors, or family members.
- Software may not be downloaded without prior approval.
- Users are not allowed to make copies of company-owned software.
- Users are required to use password-protected screen savers during the day and shut down their computers each night.
- Only IT staff members are allowed to add hard drives or install networking devices on office computers.
- The network administrator needs to assign a username and password to anyone who connects to the office network from a remote location. In addition, any remote PCs used to connect to the network must be protected with firewall and antivirus software.

Administrator duties can be covered by a specific portion of a security policy called a **privileged access policy**. This policy covers the access that administrators can have to network resources and specifies whether they are allowed to run network-scanning tools, run password-checking software, and have root or domain administrator access.

Using Security Policies to Conduct Risk Analysis

After a security policy has been drafted, the work does not end there. Employees and managers are educated about the policy's requirements and provisions, and then the safeguards are put in place. In addition to information gathered by existing security systems, such as log files, network traffic signatures, and peak traffic loads, the policy can specify further network monitoring with the installation of firewalls and IDPSs. You can use this information to determine how the network is performing and apply the data in future rounds of risk analysis.

It is up to you to decide how often to perform rounds of risk analysis. After you make needed adjustments to the security configuration in response to threats that have occurred, you might decide to conduct ongoing risk analysis every three or six months; there's no firm rule. Conducting risk analysis is important after a major change, such as when new equipment is installed or an extensive number of staff leave the company. Each risk analysis involves the same steps as the first round of risk analysis. With each subsequent analysis, you have more real-world data for evaluating risk and its consequences.

Developing a Security Policy

The following sections outline the steps of creating a security policy and offer tips and examples for the types of policies you can create. These different types of policies are combined to form a complete security policy.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Steps to Creating a Security Policy

After conducting a risk analysis, you should summarize your findings in a report that points out the most urgent risks the company needs to address. You might summarize these risks in an introductory paragraph or section, and then present a grid that shows the most important assets in the company and the level of risk you have determined for them. Table 13-5 shows an example of this grid.

Asset	Threat	Probability	Consequences	Risk assessment
Physical assets	Low	Low	Significant	Medium
Data	Medium	Medium	Damaging	High
Software	Negligible	Minor	Minor	Low
Personal assets	Low	Low	Significant	Medium
Hardware	High	Medium	Damaging	High

© Cengage Learning 2014

Table 13-5 Risk analysis example

Remember to include definitions of types of assets and to describe your assessment method. Also, explain the risk assessment you assigned to each category of assets. For example, in Table 13-5, the risk to hardware might be high because servers are left in openly accessible work areas that are not protected by locks or any special environmental controls.

After you have approval from management, follow these steps to create the security policy:

1. Form a group that meets to develop the security policy. Because of the political nature of security policies and their effect on employees, be sure to include a senior administrator, a member of the legal department, IT staff, and a representative of rank-and-file employees.
2. Determine whether the organization's overall approach to security should be restrictive or permissive. A restrictive approach limits activity on the network to a few authorized tasks. A permissive approach allows traffic to flow freely and restricts only vulnerable ports, services, or computers.
3. Identify the assets you need to protect. You should have already laid the groundwork for this step in your risk analysis.
4. Determine what information needs to be logged and which network communications need to be audited. Then decide how often the results should be reviewed. **Auditing** is the process of reviewing records of network computer activity; these records identify who is connecting to a computer, what resources are requested, and whether access is granted or blocked. This information is typically recorded in a log file.
5. List the security risks that need to be addressed.
6. Define acceptable use of the Internet, office computers, passwords, and other network resources.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

7. Define the security controls to be implemented.
8. Create the policy. The risks specified in the policy can be separate sections of one document or separate policies that together make up the organization's overall security policy.

Identifying Security Policy Categories

After the security policy group has determined what should go into the policy, the contents need to be written. Members of the group can be assigned to prepare different parts of the policy. A team approach ensures that all relevant points are covered. The following sections briefly describe common categories you are likely to need.

Acceptable Use An **acceptable use policy** establishes how company resources must be used. An acceptable use policy might state the following:

> *The following acceptable use policy covers the use of Company's computers, network components, software applications, and other hardware. The term "other hardware" includes, but is not limited to, personal computers, laptops, cellular phones, flash drives, removable disks, CD-ROM/DVD drives and disks, servers, cables, routers, and tape backup systems. The term "user" is defined as a person who has an account to use Company's network resources. All users of Company's network are expected to conduct themselves in a responsible, legal, non-threatening manner at all times, specifically:*
>
> - *Users are not allowed to make unauthorized copies of copyrighted software except with the permission of the copyright holder.*
> - *Users are responsible for storing personal data on their computers. If they need assistance with storing data, they should consult the network administrator.*
> - *Users are not permitted to engage in any activity, online or offline, that harasses, threatens, or abuses other users.*
> - *E-mail accounts are for business use only; personal e-mail messages and messages that might be judged obscene, harassing, or offensive shall not be sent from or stored on Company's systems.*
>
> *Users who violate these policies will be reported to security staff. Offenses might result in loss of network privileges or termination of employment. If the offense warrants, the company may press civil or criminal charges against the user.*
>
> *I have read and understand Company's acceptable use policy and agree to abide by it.*

The user's signature follows; some companies include the signature of a witness and the date the policy was signed. These provisions provide a few examples of what can be included in an acceptable use policy. Your own policy can go into more detail regarding the use of network resources.

The acceptable use policy is usually included at the beginning of a security policy because it affects most employees in an organization and it can generate the most controversy. Because everyday work routines might have to be changed to comply with the new policy, gaining

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

cooperation from all employees can be difficult. The transition is easier if upper management and direct supervisors begin an awareness campaign early. Also, having a representative of rank-and-file employees on the policy development group is wise for a few reasons. First, employees' interests are represented. Second, their input on decisions can help avoid making changes that hurt employee productivity. Finally, having an employee representative on the committee makes managers and supervisors more likely to talk to other employees about the policy, why it is important, and why employees should support it. These discussions are usually called a security user awareness program, and should be part of the implementation plan. Efforts should start as early as possible to get employees involved in the policy.

The first step in an awareness effort could be posting memos on bulletin boards. Then supervisors might call departmental meetings to talk to employees about the new policy and answer questions. A company can also post information about the policy on its internal Web site, create brief quizzes about the security policy, or give employees hypothetical situations to consider. A company can also create games based on the quizzes. Departments could be awarded points for correct answers, and the department with the most points might win a party or bonuses.

The point is to make sure that a policy is not sprung on unsuspecting employees. That approach would practically ensure failure because the policy depends on employees following it. An awareness program also needs to explain how the policy benefits employees. For example, the policies could lead to increased revenues from using new technologies, which might result in better pay or improved working conditions.

Extranets and Third-Party Access Allowing third-party access to the company network is often a requirement when outsourcing jobs or working with contractors. An **extranet** is a private network that a company sets up as an extension of its corporate intranet to allow contractors, suppliers, and external partners access to a limited portion of the network infrastructure. Parties who have access to the extranet should be included in the security policy. This third-party access policy should include the following points at a minimum:

- Access should be permitted only for company business.
- Third parties should be subject to a security screening process.
- Precise methods for allowing and denying connectivity should be defined.
- The duration of permitted access and the details of terminating access should be defined.
- Penalties and consequences for violating access terms should be defined because they are different from those for employees.

User Accounts, Password Protection, and Logical Access Controls By creating a policy that defines how user accounts are used, you gain flexibility in developing and enforcing the security policy because you do not limit yourself to employees who work onsite. User accounts include employees and contractors who work at home and suppliers who connect to your network from their own facilities. Your security policy might include the following specifications for user account policies:

- Users are not permitted to gain access to an unauthorized resource.
- Users cannot block an authorized user from gaining access to an authorized resource.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Users cannot give their account usernames and passwords to other people for any reason. If a password is lost or a user account is disabled, contact the administrator or help desk for assistance.
- Users must protect their usernames and passwords in a secure location that is not visible on their desktops.
- Users must abide by the password policy of the company, specifically:
  - Passwords must meet complexity requirements. Strong passwords use a random combination of letters, numbers, and symbols, and use both uppercase and lowercase characters.
  - Passwords must be at least eight characters.
  - Passwords must not be words from the dictionary, names, dates, or other information that can be associated with the user or company.
  - Passwords must be changed every 90 days.
- Users may not reuse old passwords for a period of one year.

Passwords represent a first line of defense for many organizations. They allow users to access e-mail, control access to the network from outside, control access to shared directories on servers, and more. Companies often require users to protect their computers or Web browsers with passwords so that other staff or visitors cannot use the resources after hours. In any case, the security policy should include a password policy that is enforced by system software whenever possible.

Companies must ensure that employees have ready access to the information needed to perform their jobs but that access to confidential information is limited. For example, many employees have access to read certain information but not to change or delete it. Unlimited access to this information might be granted only to a few employees. Logical access controls are software components used as a protection mechanism for limiting access to information.

Remote Access and Wireless Connections Organizations often use freelancers and consultants who work via broadband connections, e-mail, virtual private networks (VPNs), and wireless connections. In addition, mobile workers often need to connect to the office while on the road or at home, and business partners want to update their orders or view account information by connecting to the company network.

These types of workers represent an opportunity for increased productivity but they also increase security vulnerabilities. If a user connects to the corporate network from a computer that is infected with a virus or has been compromised by an attacker, the entire corporate network is threatened. As another example, mobile devices can be stolen, and the thieves can then attempt to gain access to your network.

A **remote access and wireless connection policy** specifies the use of **role-based authentication**, which gives users limited access based on their assigned role in the company and the resources and applications the role is allowed to use. The sensitivity of these applications determines the type of authentication used. For example, access to applications or data sources that are not confidential can be granted with a simple password. Access to confidential resources can be secured with a smart card or token, which is a piece of hardware used

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

with a password to provide **two-factor authentication**. This type of authentication requires a combination of identifying a physical property of the user, such as fingerprints, a signature, or a retina scan; identifying a physical item such as a card or token; or using known information such as a password or PIN. Any two of these factors combined qualify as two-factor authentication. The idea is to add an extra layer of protective authentication to double the security. For example, if you could get money from an ATM or pay a bill simply by presenting your bank card, a thief would need only the card to use your bank account. Because you also have to enter a PIN or sign a receipt, a thief has a more difficult time using your card.

VPNs and wireless connectivity provide additional options for network access; they are popular because they are effective and inexpensive. VPNs create a tunnel to transport information through public communications media, such as regular phone lines. The data is kept safe by **tunneling protocols** and encryption. Wireless connections offer the same range of tools that are available on wired connections without the burden of cabling.

However, wireless connectivity has its own security problems. Your security policy must specify who can use wireless connectivity, how it is used, and where. A relatively new phenomenon in business organizations is called BYOD (bring your own device). Because so many workers use their own digital wireless devices, such as smartphones and tablets, companies have found that they need policies to control these devices when used for business. These policies are relatively new to businesses, but university IT departments have long supported user-owned devices in classrooms, libraries, and dormitories.

Secure Use of the Internet and E-mail An **Internet use policy** can be integrated with an acceptable use policy or the overall security policy. However, because Internet use is becoming so integral to day-to-day work, it is worthwhile to create a separate section of the security policy that covers how employees can access and use the Internet.

A clear policy that governs the use of e-mail is essential to prevent problems such as mail servers being flooded by nuisance messages. A common example is a virus hoax, which states that an e-mail is circulating with a deceptive heading such as "You're a Winner!" The hoax message warns users not to read the e-mail because it will erase all files on the user's hard disk. An employee who receives such a message might forward it to all other employees, thinking he is doing them a favor. If no policy is defined for how employees should handle nuisance e-mails, the company's mail servers could become flooded and IT employees might have to diagnose the trouble, which wastes staff time and resources.

An Internet use policy prohibits broadcasting e-mail messages. Instead, users should contact the network administrator about suspicious e-mail. The policy should also specify whether users can download software or streaming media from the Internet; if downloads are allowed, limits must be set for the size of file downloads. The policy can prohibit users from opening executable e-mail attachments that might contain viruses. It could also specify whether the company has blocked objectionable Web sites and inform users how the company protects their privacy with regard to e-mail.

Network Security A **network security policy** should clearly define and establish responsibilities for using the network and for protecting information that is processed, stored, and transmitted on the network. The primary responsibility for the network can be assigned to

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

the data owner—the manager within the organization who creates and processes the data. Secondary responsibility is then assigned to users who have access to the information. A person with primary responsibility must ensure that policies and procedures are followed to protect data.

A network security policy should clearly define the roles of people who are responsible for maintaining the network's availability. This policy should cover the following points:

- *Applicability*—What constitutes the network environment and what parts, if any, are exempt from the policy
- *Evaluations*—The value of information stored on the network
- *Responsibilities*—Who is responsible for protecting information on the network
- *Commitment*—The organization's commitment to protecting information and the network

The following sample network security policy defines responsibilities for these employees:

- *Functional managers*—Employees who have primary responsibility
- *Local administrators*—Employees who are responsible for ensuring that end users have access to network resources on their servers
- *End users*—Any employees who have access to the organization's network; these employees are responsible for using the network in accordance with the network security policy

All data users are responsible for complying with the network security policy and for reporting any suspected security breaches to management.

The following is an excerpt from a sample network security policy developed by the National Institute of Standards and Technology:

> *A computer security incident is any adverse event whereby some aspect of computer security could be threatened: loss of data confidentiality, loss of data or system integrity, or disruption or denial of availability. In a LAN environment, the concept of a computer security incident can be extended to all areas of the LAN (hardware, software, data, transmissions, etc.), including the LAN itself.*

Within this sample policy, contingency plans in a LAN environment should be developed so that any LAN security incident can be handled in a timely manner, with minimal impact on the organization's ability to process and transmit data. A contingency plan should consider incident response, backup operations, and recovery.

The purpose of incident response is to mitigate the potentially serious effects of a severe LAN security problem. A response requires the capability to react to incidents and resources to alert and inform users, if necessary. Incident response requires the cooperation of all users to ensure that incidents are reported and resolved, and that future incidents are prevented.

Backup operation plans are prepared to ensure that essential tasks can be completed after the LAN environment is disrupted and continue until the LAN is sufficiently restored. These essential tasks are identified in a risk analysis.

Recovery plans are made to permit smooth, rapid restoration of the LAN environment following interruption of LAN usage. Supporting documents should be developed and

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

maintained to minimize the time required for recovery. Priority should be given to applications and services that are critical to the organization's function. Backup operation procedures should ensure that critical services and applications are available to users.

Server Security A server security policy does not usually cover rank-and-file employees. Its purpose is to regulate IT staff who have privileged access to company servers. This policy should cover all servers, including Web and database servers. At a minimum, the policy should encompass the following areas:

- Names and positions of IT staff who are responsible for operating and maintaining servers
- Specific identification for all servers, including serial numbers and part numbers
- Username and password security requirements
- Configuration details, including hardware and software versions
- Monitoring requirements and schedules as well as logging requirements
- Data and system backup requirements, storage, schedules, and responsibilities
- System audit requirements and schedules
- Policy compliance and enforcement

Physical and Facility Security Physical and facility security encompasses a broad range of issues related to locking down hardware components of a corporate network. Computer facility security must be integrated into the overall security policy for the entire corporate facility. Common sense plays a major role in designing adequate physical security and formulating a sound physical security policy:

- A separate enclosed space should be set aside to house servers and other essential computer network components.
- The computer facility should be located on the building's ground floor. If the building is in an area prone to flooding or other forms of environmental hazards, the facility should be on a higher floor.
- The facility should have no windows and a limited number of doors—preferably one door, if the fire code permits. If possible, the facility should be located away from exterior building walls.
- All access points, including ventilation shafts, should have unbreakable coverings and be double locked. Access codes should be held by as few people as possible. Not everyone who works in the facility requires access codes.
- Access should be limited to people who work in the enclosed facility. If cleaning or maintenance workers have access to the facility, they should be supervised at all times by a facility staff member.
- Fire suppression and intrusion alarm systems should be in place.
- A 24-hour video surveillance system should be maintained and viewed regularly.
- Secure off-site storage should be arranged for backup data.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Defining Incident Handling Procedures

When formulating a security policy, you should describe in detail who responds to security incidents, what needs to be done in response, and why these procedures are necessary. This portion of a security policy is called the incident response section. You should begin by describing the need for careful and expeditious handling of an intrusion if it occurs. You might also describe the types of incidents that need to be addressed, such as the following:

- Alarms sent by intrusion detection and prevention systems
- Repeated unsuccessful logon attempts
- New user accounts that suddenly appear without explanation
- New files with unfamiliar filenames that appear on system servers
- Unexplained changes to data or deletion of records
- System crashes
- Poor system performance

The incident handling process should be outlined in the security policy. Tell employees how to identify whether an intrusion has actually occurred. For example, employees can perform this check by auditing the system to see whether new files have been added. If so, they need to determine what resources have been compromised. The affected resources should then be contained, viruses or other files introduced into the system should be eradicated, and resources should be recovered.

Assembling a Response Team

If an incident occurs, the security policy should identify exactly which security staff members should be notified. Include e-mail addresses and phone numbers. Also, specify a location where team members should assemble in case they cannot access the facility or communicate. Teamwork is essential in responding to network security incidents successfully. Organizations often designate a **security incident response team (SIRT)**—a group of employees that takes countermeasures when an incident is reported.

A SIRT might also be referred to as a computer emergency response team (CERT), which can respond to any type of system failure, not just a security-related intrusion.

A SIRT responds to security-related breaches and usually includes functions such as the following in its mission statement:

- Reacting to security breaches that originate outside or inside the organization
- Isolating, reviewing, and interpreting information about security incidents
- Assessing the extent of damage caused by a security incident
- Determining the causes of intrusions and other incidents and recommending countermeasures to prevent them from recurring
- Monitoring the integrity of the organization's network on an ongoing basis

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Typically, a SIRT contains IT operations and technical support staff, IT application staff, a chief security officer, and other information security specialists. Some large organizations have a special position called incident response manager; this person is designated to respond to incidents, perform an initial assessment, and summon the SIRT and other employees as needed. A SIRT can also include members from other areas of the organization, such as department management, public relations, and legal counsel. All SIRT members should be identified in the security policy; if the primary contact person cannot be found, another person on the list can be summoned.

Specifying Escalation Procedures

An **escalation procedure** is a set of roles, responsibilities, and measures taken in response to a security incident. It describes how an organization increases its state of readiness when a threat or security incident occurs. Incidents are usually divided into three levels of escalation:

- *Level One (minor to moderate)*—These incidents are the least severe and typically must be managed within one working day of their occurrence.
- *Level Two (major)*—These incidents are moderately serious. They should be managed on the same day they occur—ideally, within four hours.
- *Level Three (catastrophic)*—These incidents are the most serious; they must be handled immediately.

Escalation procedures also specify the employees who handle each level. A large organization might have a full-fledged department whose members are assigned only to maintain security. They often have titles such as security analyst, security architect, and chief security officer. Many organizations assign technical staff to these roles; these staff members are responsible for handling escalation in addition to their other responsibilities. In either case, escalation procedures in a security policy must specify who needs to respond.

In addition, Level Two or Level Three incidents might require the participation of outside security groups, such as Computer Incident Advisory Capability (CIAC). This highly regarded organization keeps records of serious security attacks; if your organization is hit by a new virus or an unusually strong distributed denial of service (DDoS) attack, you should inform others about the attack.

If an intrusion occurs and damage is more severe than originally thought—or if the intrusion is in progress and files are being accessed—the security policy should describe stages of response that escalate along with the incident's consequences. To determine how a response might escalate, you can develop a system for rating an incident's severity using the rankings listed in Table 13-6. Each ranking could then be mapped to an escalation chain—a hierarchy of employees who should respond to incidents and make decisions. Table 13-6 shows one possible mapping.

Table 13-6 could be extended to list actions to take in response to incidents, based on their severity. Incidents of minor to moderate severity might require a virus scan or log file review; major incidents might call for disconnecting the local network from the Internet or other network segments during the response.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Incident severity	Escalation chain
Catastrophic (Level Three)	Business owner or manager, chief security officer, senior network administration staff, all department heads, public relations officer
Major (Level Two)	On-duty manager, on-duty network administration staff, security architect
Moderate (Level One)	Immediate supervisor, on-duty security analyst, help desk
Minor or insignificant	Immediate supervisor

© Cengage Learning 2014

Table 13-6 Mapping an escalation chain to incident severity rankings

Responding to Security Incidents

To determine how incidents should be escalated, the security policy's section on incident handling should clearly define incident types and the level of escalation each one represents. Note the following examples:

- *Loss of password (Level One incident)*—The on-duty supervisor should be notified within 24 hours. The supervisor determines whether a change of password is necessary.
- *Burglary or other illegal building access (Level Two incident)*—If an unauthorized person is discovered on the premises, notify your immediate supervisor, who determines whether police need to be notified and whether the incident should be escalated to the next person on the chain. If the incident is serious enough, the supervisor notifies the appropriate people in the escalation chain, such as the CSO, IT manager, or security administrator (SA). The unauthorized person should be escorted from the building, either by the proper person in the chain or by police. The responsible party specified in the security policy then writes an incident report.
- *Property loss or theft (Level Two or Level Three incident)*—If company property has been stolen, the human resources director or the on-duty supervisor should be notified immediately. This official escalates the incident to the CSO and local law enforcement, if needed.

Including Worst-Case Scenarios

Worst-case scenarios are descriptions of the worst consequences that can befall an organization if a threat happens. These scenarios might be unlikely, but they can help you determine the value of a resource at risk. Values are derived from reasonable consequences of files, computers, and databases being unavailable for specified periods. You might prepare scenarios that account for several periods of time ranging from a few minutes to several months.

Another way to quantify the impact of financial loss or interruption of business activities is to assign a numeric value to an asset based on a range of dollar amounts. If the loss is estimated at $100 or less, for example, a 1 might be assigned to the loss; if the loss is between $100 and $1000, a 2 is assigned, and so on.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Updating the Security Policy

You should update the security policy based on incidents reported as a result of ongoing security monitoring and any new risks your company faces. Any changes to the policy should then be made available to all employees, either by e-mail or by posting the changes on the company's Web site or intranet.

The goal of changing the security policy is to change employees' habits so that they behave more responsibly. Ultimately, a security policy should result in physical changes to the organization's security configuration. For example, a call for redundant systems in the security policy might result in major expenditures for a new firewall or server that acts as a failover device. The need to review security logs daily, as prescribed in a security policy, might mean investing in log file analysis software to make the reviews easier. Better protection means fewer internal or external incidents, which enables the company to focus on its primary mission.

Conducting Routine Security Reviews

After a security policy has been in place for a while, the security policy group should decide how to conduct a routine reassessment of risks to the company and its assets. You might not need to examine every asset in the company during each reassessment; for example, you might decide to focus on the most urgent security risks or new risks that have emerged because of changes in the company.

When reevaluating the organization's security policy, keep the following requirements in mind:

- The reviews need to be routine.
- Upper management must authorize the reassessment schedule.
- The organization needs to respond to security incidents as they occur.
- The organization needs to revise the security policy because of incidents and other identified risks.

An effective security policy describes immediate steps to take when an intrusion is detected and specifies how often risk analyses should be conducted. A section of the policy should identify the people who conduct risk analysis and then describe the circumstances in which a new risk analysis is required. For instance, when new equipment is purchased, the risk analysis group should determine whether new statements need to be written and added to the security policy, and whether new security measures are needed.

Even though the security policy might specify a risk analysis every six or 12 months, the policy should be flexible enough to allow emergency reassessments as needed. For instance, any attacks on partner businesses or remote offices should prompt the security policy group to reassess risks for the organization. In addition, news of major security attacks on Internet servers or viruses circulating on the Internet should prompt risk reassessment.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Chapter Summary

- Risk analysis plays a central role in formulating a security policy. Risks should be calculated and security policies should be amended on an ongoing basis as the network configuration changes and new threats emerge.
- Risk analysis covers a company's computer hardware, software, and informational assets. It should be performed before *and* after a security policy is created. The goal is not to reduce risks to zero but to manage risk at reasonable levels on an ongoing basis.
- The first task during risk analysis is to identify the assets you need to protect. Then you assess threats to your network, such as attacks, power outages, and environmental disasters. Next, you determine the probability that the threats might occur.
- After assessing the level of risk to assets, you need to determine countermeasures and safeguards for minimizing risk. Decide how to secure computing assets; logical assets such as IM records, e-mail, and log file records; data stored in your databases; application software; and personal assets of employees. You then develop a plan for conducting routine risk analyses and a plan for handling security incidents.
- To perform a risk analysis, you can use an approach such as Survivable Network Analysis (SNA) or Threat and Risk Assessment (TRA). A risk analysis describes the level of risk to each asset in the organization as well as the economic impact if the asset is lost or damaged.
- A security policy provides a foundation for an organization's overall security stance. The policy gives employees guidelines for handling sensitive information and provides the IT staff with instructions for configuring defensive systems. A security policy also reduces the risk of legal liability for a company and its employees.
- To protect overall security, it is important to formulate a clear policy that explains employees' rights and how they should handle company resources.
- Because risk cannot be eliminated, businesses sometimes purchase cyber risk insurance to offset potential losses. A security policy can help an organization apply for this insurance by identifying the assets that need to be insured and threats that require protection.
- A good security policy is based on risk assessment, covers acceptable use of system resources, sets priorities for critical resources that need to be protected, and specifies the use of network resources by administrators and security staff.
- Legal liabilities should be covered in a security policy. Statements of acceptable and unacceptable use must be included, and guidelines for violations and punishments must be covered. This information helps the company avoid litigation from employees who are fired for misuse of systems.
- If a security incident is caused by a criminal offense, it is important to understand your legal obligations and how to protect yourself from litigation. After an investigation is turned over to the police, Fourth Amendment protections for search and seizure apply; failure to respect these protections can result in lawsuits. You must understand the implications of a criminal investigation and comply with court-ordered subpoenas and search warrants.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- A security policy is often formulated as a series of specific policies rather than one long document. Creating a security policy requires forming a group to develop the policy; determining the overall security approach; identifying assets to protect; specifying auditing procedures; listing security risks; defining acceptable use; and creating the specific policies, such as those for user accounts, password protection, Internet use, and remote access.
- You need to present a proposal to management and gain approval before you create the security policy. This process involves explaining the expected ROI and calling attention to other costs associated with security incidents, such as loss of productivity, loss of morale or customer confidence, and possible legal costs.
- A security policy should describe in detail who responds to security incidents, what needs to be done in response, and why the procedures are necessary. Collectively, this information is called the incident response section of the security policy. It defines the response and escalation for incidents of varying severity and contains contact information for people who need to be notified. Organizations often have a security incident response team (SIRT) to take countermeasures when an incident is reported.
- An escalation procedure should be defined to determine who is notified during each type of incident. You can define different incident levels, such as minor, moderate, major, and catastrophic. As the level of severity increases, higher-level managers in the organization must respond.
- Security policies should be reviewed and updated regularly. They should also be modified to account for intrusion attempts, actual intrusions, personnel changes, and major equipment purchases.

Key Terms

acceptable use policy A policy that defines acceptable and unacceptable uses of company resources.

assets The hardware, software, and informational resources you need to protect by developing and implementing a comprehensive security policy.

auditing The process of reviewing records of network computer activity; these records identify who is connecting to a computer, the resources being requested, and whether access is granted or blocked.

cost-benefit analysis A technique for comparing the costs of an investment with the benefits it proposes to return.

electronic assets Word-processing documents, spreadsheets, Web pages, and other documents on your network computers.

escalation procedure A set of roles, responsibilities, and measures taken in response to a security incident.

exposure Vulnerability to loss resulting from the occurrence of a threat, such as disclosure, destruction, or modification of information resources. Exposure increases with the presence of multiple threats.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

extranet A private network that a company sets up as an extension of its corporate intranet so contractors, suppliers, and external partners can access a limited portion of the network infrastructure.

fault tolerance The ability of an object or a system to continue operating despite a failure.

Fourth Amendment An amendment to the U.S. Constitution that provides protection from illegal search and seizure.

Internet use policy A policy that defines how users can access and use the Internet and that specifies what rules apply to e-mail and other communications, such as instant messaging.

Monte Carlo simulation An analytical method that simulates a real-life system by randomly generating values for variables.

network assets Routers, cables, bastion hosts, servers, and firewall components that enable employees to communicate with one another and other computers on the Internet.

network security policy A policy that defines and establishes responsibilities for protecting the network and the information processed, stored, and transmitted on the network.

privileged access policy A policy that details additional access options and responsibilities of users with privileged access to resources.

probability The possibility that a threat will actually occur, as influenced by geographic, physical, habitual, or other factors that increase or decrease the likelihood of occurrence.

remote access and wireless connection policy A policy that defines necessary security measures before a remote desktop or wireless connection is added to an organization's network.

residual risk The risk remaining after countermeasures and defenses are implemented.

risk The possibility of incurring damage or loss.

risk analysis A process of analyzing the threats an organization faces, determining which resources are at risk, and determining the priority of each asset.

risk management The process of identifying, choosing, and setting up countermeasures justified by the risks you identify.

role-based authentication A method of authentication that grants users limited system access based on their assigned role in the company and that defines the resources the role is allowed to use.

search warrant A legal document issued by a court that allows authorities to search a particular place for specific evidence. The warrant must detail what the search is seeking and where authorities are permitted to look for it.

security incident response team (SIRT) A group of people designated to take countermeasures when an incident is reported.

security policy A statement that defines the defenses configured to block unauthorized access to a network, acceptable use of network resources, an organization's response to attacks, and how employees should handle the organization's resources to prevent data loss or file damage.

subpoena A legal document that requires a person to appear in court, provide testimony, or cooperate with law enforcement.

survivability The ability to continue functioning in the presence of attacks or disasters.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Survivable Network Analysis (SNA) A security process that assumes a computer system will be attacked and follows a set of steps to build a system that can survive the attack.

Threat and Risk Assessment (TRA) An approach to risk analysis from the standpoint of threats and risks to an organization's assets and the consequences if those threats and risks occur.

threats Events and conditions that could potentially occur; their presence increases risk.

tunneling protocols Network protocols that encapsulate (wrap) one protocol or session inside another.

two-factor authentication Authentication that requires at least two forms of verification for a user to be granted access.

vulnerabilities Situations or conditions that increase threats, which in turn increase risk.

worst-case scenarios Descriptions of the worst consequences that can befall an organization if a threat occurs.

Review Questions

1. Survivable Network Analysis begins with what assumption?
   a. that you have laid the groundwork for a risk analysis
   b. that your network will be attacked
   c. that the probability of threats is increasing constantly
   d. that an effective security policy can reduce risks to zero

2. To determine the value of hardware and software you need to protect, which of the following approaches is easiest to use?
   a. getting the most recent prices online
   b. keeping records of purchase costs
   c. using your experience and expertise
   d. interviewing support personnel

3. When should an organization conduct a new round of risk analysis?
   a. every month
   b. every three months
   c. as frequently as possible
   d. when equipment or staff change significantly

4. A risk analysis report should call attention to _____.
   a. all identified risks
   b. the most urgent risks
   c. the newest risks
   d. the risks that are easiest to manage

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

5. The ultimate goal of a security policy is which of the following?
   a. reducing the risks to zero
   b. doing it right the first time so the policy does not have to be rewritten constantly
   c. convincing management that the IT budget should be increased
   d. none of the above

6. What are the hardware, software, and informational resources you need to protect?
   a. threats
   b. tangibles
   c. assets
   d. business holdings

7. Ensuring that databases remain accessible if primary systems go offline is known as _____.
   a. fault tolerance
   b. failover
   c. redundancy
   d. resiliency

8. Which of the following technologies helps protect sensitive data even after it has been stolen from a secured medium?
   a. virus protection
   b. authentication
   c. encryption
   d. Spybot

9. Which of the following sections of a security policy affects the most people in an organization?
   a. incident handling policy
   b. privileged access policy
   c. acceptable use policy
   d. remote access policy

10. What is an escalation procedure? (Choose all that apply.)
    a. It describes how network security can be improved in stages.
    b. It describes how a virus can multiply and affect more assets.
    c. It describes different levels of response based on incident severity.
    d. It identifies employees who should be involved in the response.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

11. Which of the following, if worded correctly, can protect companies from wrongful termination lawsuits?
    a. nondisclosure clauses
    b. acceptable use policies
    c. penalty clauses
    d. punitive clauses

12. A password policy might specify which of the following attributes for password selection?
    a. length requirements
    b. complexity requirements
    c. frequency for changing passwords
    d. all of the above

13. Which of the following provides employees with formal instructions about the organization's security strategy?
    a. acceptable use policy
    b. risk assessment
    c. strategy meeting
    d. security user awareness program

14. If organizations have employees who connect remotely, which of the following security concerns should be considered?
    a. the possibility of mobile devices being stolen
    b. virus infections spreading from home and mobile systems to corporate systems
    c. the use of updated, effective antivirus and firewall software on mobile devices or home systems that connect to the network
    d. all of the above

15. A password policy should be established in the _____ and enforced by _____ whenever possible.
    a. risk assessment process, management
    b. company Web site, network administrators
    c. security policy, software
    d. company employee handbook, security guards

Hands-On Projects

Hands-On Project 13-1: Calculating Replacement Costs

Time Required: 15 minutes

Objective: Use a risk analysis tool to calculate replacement costs for equipment.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Description: A software tool can help you perform risk analysis with consistency and objectivity. In this project, you download a trial version of a Windows tool called Project Risk Analysis (PRA) by Katmar Software. You then enter information about network resources in your school's lab and calculate the contingency funds needed to replace lab equipment if disaster strikes. You will need a file-archiving utility, such as WinZip. Enter values for computer equipment in your lab and estimate how much data is contained on each piece of equipment. You may need to research the cost of the network resources. The project specifies the names and costs of sample devices, but you can substitute the names and costs of your lab equipment. If you are not in your lab, you can use the information shown in the project steps.

1. Log on to either Windows Server 2008 or Windows 7.
2. Start your Web browser and go to **www.katmarsoftware.com/pra.htm**.
3. Read the description of the program, and then click **Download Now!**.
4. When you are prompted to open or save the Projrisk.zip file, click **Save** and then save the file to a folder in your file system.
5. When the download is finished, click **Close** in the Download Complete window. Exit your Web browser, and then double-click the file to open it with WinZip. Extract the files to the folder where you saved the .zip file you downloaded.
6. Double-click the **ProjRisk_Setup.exe** file, and follow the steps in the setup program to install the file on your computer.
7. Click **Start**, point to **All Programs**, point to **Risk Analysis**, and click **ProjRisk** to start the program. The first time you run the program, you see a window that states the terms under which the program can be run. (It runs 30 times as an evaluation.) Click **OK**, and then click **OK** again when you see a second shareware reminder.
8. Click **Add** to open the Add a Record at the End window. In the Description text box, type **Computers**.
9. In the Distribution section, click the **Normal** option button.
10. In the Likely Cost text box, type **15000**. Do not use commas when entering numbers. In the Low Cost text box, type **10000**.
11. Click **OK** to return to the main Project Risk Analysis window. Your estimate is entered in the first row.
12. Repeat Steps 8 through 11. Remember to click **Normal** in Step 9 each time. If you are not entering your own information, enter the following descriptions and costs:

Description: **Software**	Likely Cost: **5000**	Low Cost: **3500**
Description: **Printers**	Likely Cost: **500**	Low Cost: **400**
Description: **Switches**	Likely Cost: **150**	Low Cost: **100**
Description: **Cables**	Likely Cost: **100**	Low Cost: **75**
Description: **Monitors**	Likely Cost: **5000**	Low Cost: **4000**

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

13. In the main Project Risk Analysis window, click the **Analyze** button to see the Overall Cost Distribution graph. What are the Lowest Cost, Highest Cost, and Mean Cost?
14. Click the **Statistics** button. What is the mean cost in the Simulation Statistics Report window?
15. Close the Simulation Statistics Report window, and exit Project Risk Analysis. Leave your system running for the next project.

Hands-On Project 13-2: Conducting Security Policy Analysis

Time Required: 20 minutes

Objective: Evaluate security policy clauses, identify deficiencies, and update policies in response to events or changes.

Description: Security policies should be revised to address security breaches or new threats. In this project, you evaluate the theft of proprietary information and identify some obvious deficiencies in a security policy. Then you recommend changes to the security policy to prevent similar incidents from recurring.

A local branch office of a major national stock brokerage had no policy that required the termination of user ID and password privileges after employees leave. A senior trader left the brokerage and was hired by a competing brokerage. Shortly thereafter, the first brokerage lost two clients who said they were moving to a competing firm; their personal data files disappeared mysteriously from the company's databases. In addition, a year-end recommendations report that the senior trader had been preparing was released two weeks early by the competing brokerage. An investigation of the company's access logs revealed that the employee records file had been accessed by someone outside the company. The job records, however, did not reveal whether the report had been stolen because they had not been set up to record object accesses in a log.

The existing security policy states the following:

"On termination, employees shall surrender any laptops, disks, or computer manuals they have in their possession. They are no longer authorized to access the network, and they shall not take any hardware or software when they leave the office."

1. What changes would you make to the existing security policy so that security is improved after employees are terminated?
2. Brainstorm for ideas to develop a security policy clause that covers access of company records and helps track when files are accessed.

Hands-On Project 13-3: Collecting a Hardware and Software Inventory

Time Required: 15 minutes

Objective: Create a hardware and software inventory of your network and save it to a file.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Description: Creating a complete hardware or software inventory of your network can be a daunting task. Several methods, from simple VBScripts to high-end software programs, can automate this chore. In this project, you use Network Asset Tracker from MIS Utilities to gather hardware and software data about your computer and save it to a file.

1. Start your Web browser. Go to **www.misutilities.com/network-asset-tracker/index.html**.
2. Click the **DOWNLOAD** tab. Save the natracker.exe file to your desktop.
3. Right-click **natracker.exe**, and click **Run as Administrator**. In the User Account Control window, click **Yes**.
4. In the Language Selection window, click **Next**.
5. In the welcome window, click **Next**.
6. In the Readme Information window, click **Next**.
7. In the License Agreement window, click **Yes, I agree with all the terms of this license agreement**, and click **Next**.
8. In the Destination Folder window, click **Next**.
9. In the Select Components window, click **Next**. In the Program Group window, click **Next**.
10. In the Ready to Install the Program window, click **Next**.
11. In the Installation Complete window, click **Finish**.
12. Click **Start**, click **All Programs**, click **Network Asset Tracker**, and click **Network Asset Tracker** to start the program. If a User Account Control window appears, click **Yes**. Click **Close** to dismiss the Tip of the Day.
13. Right-click the **LOCALHOST** computer icon, and click **Get/Update Info**. A window like the one shown in Figure 13-8 displays details about your system.
14. Click the **Report** menu, and click **System Information**. Under Type of report, click the **List** option button, and click **Ok**.
15. Click **Export** and save the System Information file to your desktop. Click **OK** in the Information window, and click **Close**.
16. From your desktop, double-click the **System Information.html** file. Read through the file to examine the information that was gathered. When you are finished, close the file.
17. Click the **Report** menu, and click **Software Information**. Click **Ok**.
18. Click **Export** and save the Software Information file to your desktop. Click **OK** in the Information window, and click **Close**.
19. From your desktop, double-click the **Software Information.html** file. Read through the file to examine the information that was gathered. When you are finished, close the file.
20. Click the **File** menu, and then click **Exit** to close the program.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

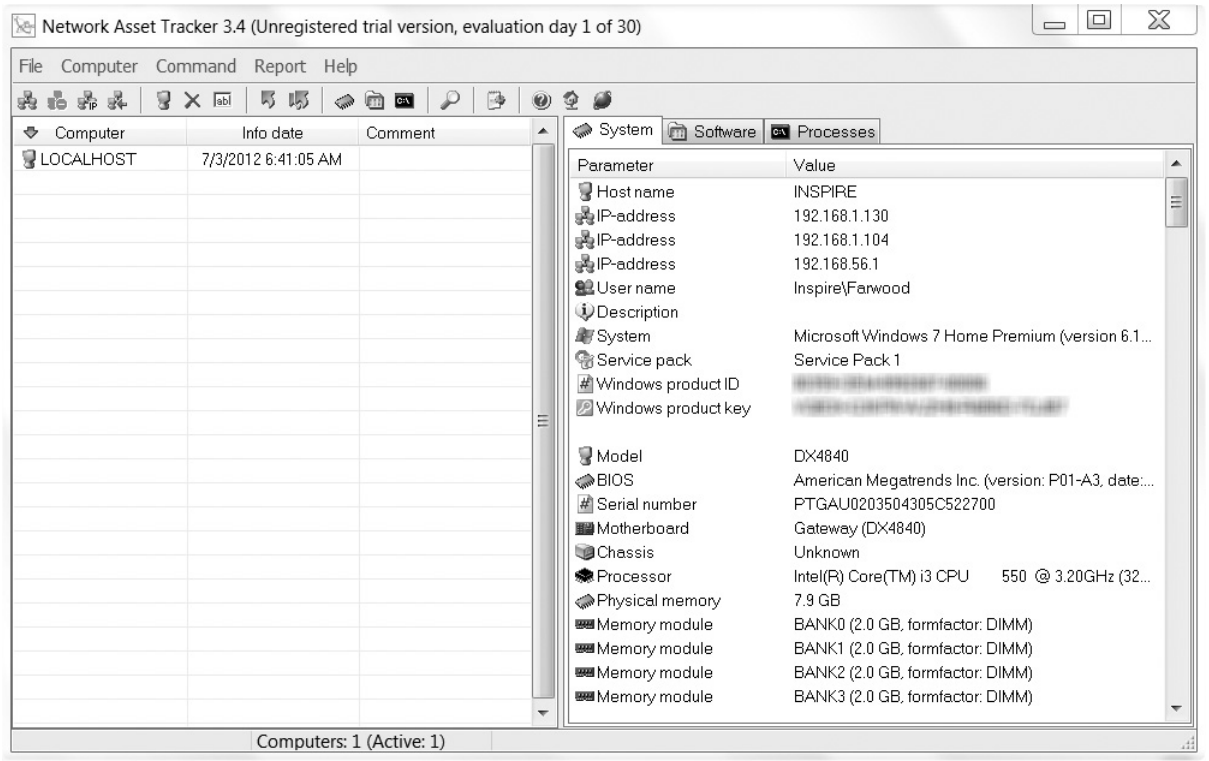

Source: MIS Utilities

Figure 13-8 Network Asset Tracker local host details

Hands-On Project 13-4: Conducting Asset Identification, Threat Analysis, and Security Controls

Time Required: 45 minutes

Objective: Learn to apply risk assessment concepts.

Description: A first step in risk analysis is identifying the assets you need to protect. In this project, you identify assets and perform a threat analysis of your environment. Your instructor will set the boundaries of the environment. Your task is to conduct a thorough, objective assessment of assets in your classroom, and then to prioritize the assets and propose ideas for safeguards. This activity can be done individually or in groups.

1. Create a chart with four columns. Label the columns **Asset**, **Threats**, **Probabilities**, and **Security Controls**. Be sure to leave enough room for your notes.
2. Starting at the classroom door, begin writing down assets. Write a short description of each asset in the Asset column, such as "student computer."
3. After you have listed all the assets in your classroom, assess the threats. Starting with the first asset listed, write down threats that could damage or destroy each asset. Examples of threats include a leaky roof, exposed wires, and food or drink that could

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

spill. Be sure to look in all directions, including the ceiling. Be sure to note drop ceilings or false floors because they are a potential point of access for intruders.

4. Next, examine the probabilities of a threat occurring. Use the information in Table 13-2 as a guide. For example, if you listed a leaky roof as a threat and you live in an area with frequent rainfall, the probability of that threat occurring and causing damage is high. If you live in an arid desert climate, that probability is not as high.
5. Next, determine security controls to reduce threats. Beginning with your first asset, look for ways to manage risks to it, and write down safeguards. If the roof is leaking above an expensive computer, the cost of repairing the roof might be warranted, but simply moving the computer reduces the risk and is more cost effective. On the other hand, if the power supply to the server room is prone to fluctuations that could damage delicate electronics, investing in uninterruptible power supplies (UPSs) or a generator could be less expensive than replacing damaged equipment. Remember that the priority level determines the investment in security.
6. If time permits, share your findings with the class. Discuss the results and justify your assessments and recommendations. Were the results similar? Discuss similarities and differences in findings.

Case Projects

Case Project 13-1: Conducting Risk Assessment and Analysis

Risk assessment can be as simple as noting an unlocked door or a password written on a slip of paper, or it can be a complicated process that requires several team members and months to complete. A large enterprise environment probably has multiple locations, diverse activities, and a wide array of resources to evaluate. You do not need such a complicated network, however, to work through this case project. The main idea is to learn how to apply your knowledge in a methodical fashion to produce useful and accurate results. Approaching a task such as risk assessment without a strategy in place usually results in repeated steps, wasted resources, and mediocre results at best. Even worse, you might miss critical information.

In a real risk analysis, one of the first steps is to meet with all department managers, upper management, employee representatives, workers in the production environment, human resources staff, and other staff members to get their input. Such a meeting is not possible in this situation, so direct any questions to your instructor or do independent research to find your answers.

In this project, you conduct a risk assessment and analysis of a small e-commerce business. You decide what product or service the business sells over the Internet. Use the following files provided by your instructor:

- Facilities diagram
- Network diagram
- Asset identification worksheet

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Business process identification worksheet
- Threat identification and assessment worksheet
- Threat mitigation worksheet

1. Give your company a name.

   Company name: _____

   Products or services: _____

2. Identify the business processes that must continue for the organization to keep functioning—for example, collecting money from customers, receiving and processing sales, and developing new products. List major business processes that drive your company in the Business Process column of the business process identification worksheet. (Use your imagination and common sense to complete this step.) Assign a priority level to each process using the priority rankings in the following list. Write down the department that performs the process. Leave the Assets Used column blank for now.
   - *Critical*—Absolutely necessary for business operations to continue. Loss of a critical process halts business activities.
   - *Necessary*—Contributes to smooth, efficient operations. Loss of a necessary process doesn't halt business operations but degrades working conditions, slows production, or contributes to errors.
   - *Desirable*—Contributes to enhanced performance and productivity and helps create a more comfortable working environment, but loss of a desirable process doesn't halt or hurt operations.
3. Identify the organization's assets. Using the asset identification worksheet, list each asset, its location, and its approximate value, if known. For multiple identical assets, describe the asset and list the quantity instead of listing each asset. In organization-wide risk assessments, you would list all assets, including office furniture, industrial equipment, personnel, and other assets. For this project, stick to IT assets, such as computers, servers, and networking equipment. List all the equipment needed to build your network as well as any cabling in the facility. Assume that the facility is already wired for a computer network and has network drops for each computer. Remember to list items such as electricity and your Internet connection.
4. Next, determine which assets support each business process. On your business process identification worksheet, list the assets needed for each business process in the Assets Used column.
5. Document each process and assign a priority to it. Next, transfer the priority rankings to your asset identification worksheet. Now you know which assets are the most critical to restore and warrant the most expense and effort to secure. You also have the documentation to justify your security actions for each item.
6. Next, assess existing threats. Table 13-7 shows examples of evaluating types of threats and suggests ways to quantify them. On the threat identification and

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

assessment worksheet, list each possible threat. Be sure to consider threats from geographic and physical factors, personnel, malicious attacks or sabotage, and accidents. Also, examine the facilities diagram for flaws in the facility layout or structure that could pose a threat, such as an air-conditioning failure or loss of electrical service. For each threat, assess the probability of occurrence (POC) on a scale of 1 to 10, where 10 represents the highest probability. Enter the ratings in the POC column for each threat.

Type of threat	How to quantify
Severe rainstorm, tornado, hurricane, earthquake, wilderness fire, or flood	Collect data on frequency, severity, and proximity to facilities. Evaluate the past quality and speed of local and regional emergency response systems to determine whether they helped to minimize loss.
Train derailment, motor vehicle accident, toxic air pollution caused by accident, or plane crash	Collect data on the proximity of railroads, highways, and airports to facilities. Evaluate the construction quality of transportation systems and the rate of serious accidents on each system.
Building explosion or fire	Collect data on the frequency and severity of past incidents. Evaluate local emergency response to determine its effectiveness.
Militant group attacking facilities, riots, or civil unrest	Collect data on the political stability of the region where facilities are located. Compile and evaluate a list of groups that might have specific political or social issues with the organization.
Computer hacking (external) or computer fraud (internal)	Examine data on the frequency and severity of past incidents. Evaluate the effectiveness of computer security measures.

© Cengage Learning 2014

Table 13-7 Threat evaluation and quantification methods

7. Using the asset identification worksheet, determine which assets would be affected by each threat. List those assets in the Assets Affected column of the threat identification and assessment worksheet. For an electrical outage, for example, list all assets that require electricity to operate. For a hardware failure, list all assets that a hardware failure would disrupt, damage, or destroy.
8. In the Consequence column, enter the consequences of the threat occurring. Use the following designations:
   - *Catastrophic (C)*—Total loss of business processes or functions for one week or more; potential complete failure of business
   - *Severe (S)*—Business would be unable to continue functioning for 24 to 48 hours; loss of revenue, damage to reputation or confidence, reduction of productivity, complete loss of critical data or systems

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- *Moderate (M)*—Business could continue after an interruption of no more than four hours; some loss of productivity and damage or destruction of important information or systems
- *Insignificant (I)*—Business could continue functioning without interruption; some cost incurred for repairs or recovery; minor equipment or facility damage; minor productivity loss and little or no loss of important data

9. Rate the severity of each threat in the Severity column. Use the same designations as in Step 8 (C, S, M, or I). You derive these ratings by combining the probability of occurrence, the asset's priority ranking, and the potential consequences of a threat occurring. For example, if an asset has a Critical (C) priority ranking and a Catastrophic (C) consequence rating, it has a Catastrophic (C) severity rating. If you have mixed or contradictory ratings, you need to reevaluate the asset and use common sense to determine the severity rating. A terrorist attack that destroys the facility might have a POC of 1, depending on your location, but the consequences would definitely be catastrophic. Even so, because of the low POC, you wouldn't necessarily rank its severity as catastrophic.
10. On the threat mitigation worksheet, list assets that are ranked as the most critical and are threatened with the highest severity. In the Mitigation Techniques column, list recommendations for mitigating threats to those assets. For example, to mitigate the threat of an electrical outage damaging a critical server, you might suggest a high-end UPS.
11. Review your work, and submit it to your instructor.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

chapter 14

Ongoing Security Management

After reading this chapter and completing the exercises, you will be able to:

- Strengthen network control by managing security events
- Improve analysis by auditing network security procedures
- Strengthen detection by managing an intrusion detection and prevention system
- Improve network defense by changing a defense-in-depth configuration
- Strengthen network performance by keeping pace with changing needs
- Increase your knowledge by keeping on top of industry trends

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

In previous chapters, you learned about designing and configuring an intrusion detection and prevention system (IDPS). This chapter discusses how to manage the information from an IDPS and other security systems in your network defense. Through efficient management, you can maximize the performance of your IDPS and other areas of your network's overall security configuration.

Many organizations create positions, such as computer security manager (CSM), to oversee the ongoing management of network security systems. The responsibilities of a CSM should be spelled out clearly in the organization's security policy. In this chapter, you learn about these managerial responsibilities, including security event management, security auditing, managing an IDPS, improving a defense-in-depth configuration, keeping pace with network needs, and maintaining your knowledge to keep up with fast-changing events in network security.

Strengthening Control: Security Event Management

If you are hired by a large corporation to manage security, your initial tasks will probably include cataloging the various security devices you need to manage. You might discover that the organization already has a range of devices, such as the following:

- Packet-filtering routers
- VPN appliances
- An IDPS at each branch office
- One or more firewalls at each office
- Event logs or syslogs (system logs) for selected systems in each office

The flow of information from these devices is shown in Figure 14-1.

You could easily have 10, 20, or more devices sending log file data that you need to manage. How do you keep track of all the information and respond to legitimate security concerns without becoming overwhelmed? You need to create and manage a security incident response team (SIRT), as described in Chapter 13. You also need to establish a **security event management program** for gathering and consolidating events from multiple sources so that you can analyze the information to improve network security. The following sections discuss different aspects of this program to improve procedures and strengthen security while maintaining productivity:

- Monitoring events
- Managing data from multiple sensors
- Managing IDPS signatures
- Managing change

Monitoring Events

Network protection is needed on an ongoing basis to keep up with new threats and vulnerabilities. You need to establish a process for securing, monitoring, testing, and continually strengthening your network's level of protection. One way to improve network defenses is through ongoing **event monitoring**—reviewing alert and event logs that your security devices

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

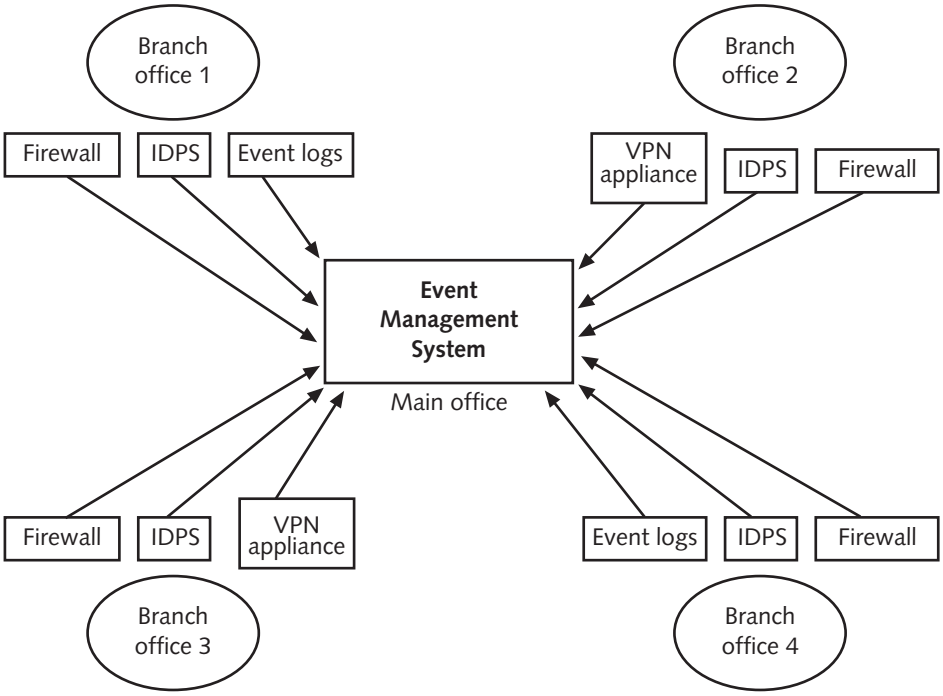

Figure 14-1 Flow of information from multiple sources

and operating systems produce and testing the network periodically to identify weak points. The goal is to strengthen defenses by gathering information, changing procedures, and improving the network. You need to monitor the following events:

- Logins
- Creation of user accounts and groups
- Correct handling of e-mail attachments
- Backups and other ways to maintain and protect sensitive information
- Antivirus scanning and control
- Procedures for granting secure remote access to mobile users

For example, suppose that your network's database server contacts a remote host via the Internet. The IDPS sends an alarm because you have configured notifications for these types of events; the database server should receive only incoming connections, not initiate its own outgoing connections. The port that the database server uses to connect to the remote host, 40449, indicates that the connection is suspicious because the port is normally closed.

To verify this connection, you can run the Netstat utility by opening a command prompt and entering the netstat –a command on the database server. Netstat reviews the current connections and reports that the server is listening for connections on the expected ports, such as port 1028 for NetBIOS connections. Netstat also reports that the server is listening for

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

connections on port 40449, which indicates that a Trojan program might have been installed on the database server and is initiating a connection to the remote host over this port. If so, an attacker might have been able to log in to the database server with a legitimate password that he "cracked" or obtained through other fraudulent means. Your response needs to incorporate several steps:

- Block the connection.
- Locate and eradicate the Trojan program.
- Determine how the intrusion occurred.
- Change network passwords.
- Possibly change the way users safeguard their passwords.

To be effective, your response must be as quick as possible. An information attack does not take long to execute, and damage can occur in an instant. You need to develop a team approach to network security and make use of automated responses, such as the alarm systems that are built into an IDPS. You also need to coordinate data from multiple sources and be aware of new network security threats.

Managing Data from Multiple Sensors

Consider an organization that has offices in several countries or a national organization with branch offices in separate states. Each office network has its own firewall and IDPS. Each IDPS has sensors that gather data passing through the gateway. Some of that data contains alert messages that need to be reviewed in a timely fashion. How do you process all the events at these offices? You need to install database software that sorts through the events and enables you to view them systematically. As a security manager, you need to address the following questions:

- Should all the data be consolidated and flow through a central security location?
- Should the data from sensors go to security managers at each office?

These questions are addressed in more detail in the following sections.

Centralized Data Collection

In **centralized data collection**, an organization's event and security data is funneled to a centralized management console in the main office. Depending on the number of branch networks and security devices your system has, an extraordinary amount of data can traverse network gateways and pass through the firewall at the main office.

When you set up a firewall or IDPS, you are usually required to identify the location of a host computer running a management console application. Data from security devices in your organization's network and its remote networks can be transmitted to that console via its IP address. Figure 14-2 shows this arrangement, in which event data from separate offices is sent to a management console in the organization's main office.

Centralized data collection offers the following benefits:

- Reduced costs because you have fewer systems to maintain
- Less administrative time
- Improved efficiency

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

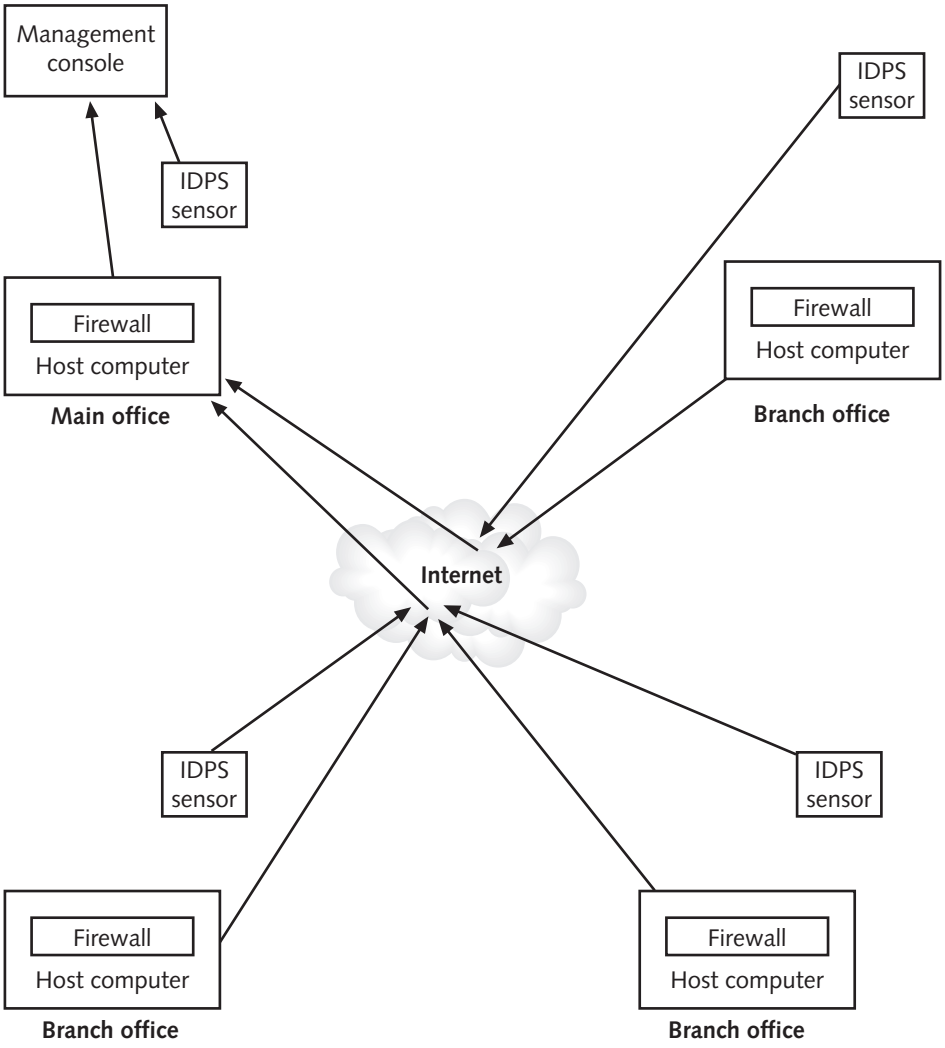

© Cengage Learning 2014

Figure 14-2 Centralized data collection

A disadvantage to this approach is that you must find a way to transmit data from each collection point to the centralized management console. The traffic must also be protected by encryption and encapsulation; a VPN connection is ideal for such a situation, as you learned in Chapter 11. However, most security devices can communicate with one another using their own encrypted "handshake" mechanism, and many companies rely on this feature instead of setting up a VPN. Unfortunately, most IDPSs and other security devices use standard ports by default that attackers can exploit.

Distributed Data Collection The second option when managing data from multiple sensors is to use distributed data collection, which reduces the amount of network traffic (see Figure 14-3). In **distributed data collection**, data from security devices such as firewalls

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

and IDPSs goes to a management console on its own local network. Security managers in each network review the data separately and then analyze and respond to events as needed. Distributed data collection requires the organization to maintain separate security managers in each branch office as well as separate management console software. This arrangement saves bandwidth, but offices must still communicate with one another about security incidents. For example, if an event at the Miami office might affect other offices, the security manager in Miami should notify other branch managers to help prevent the event from occurring in their offices.

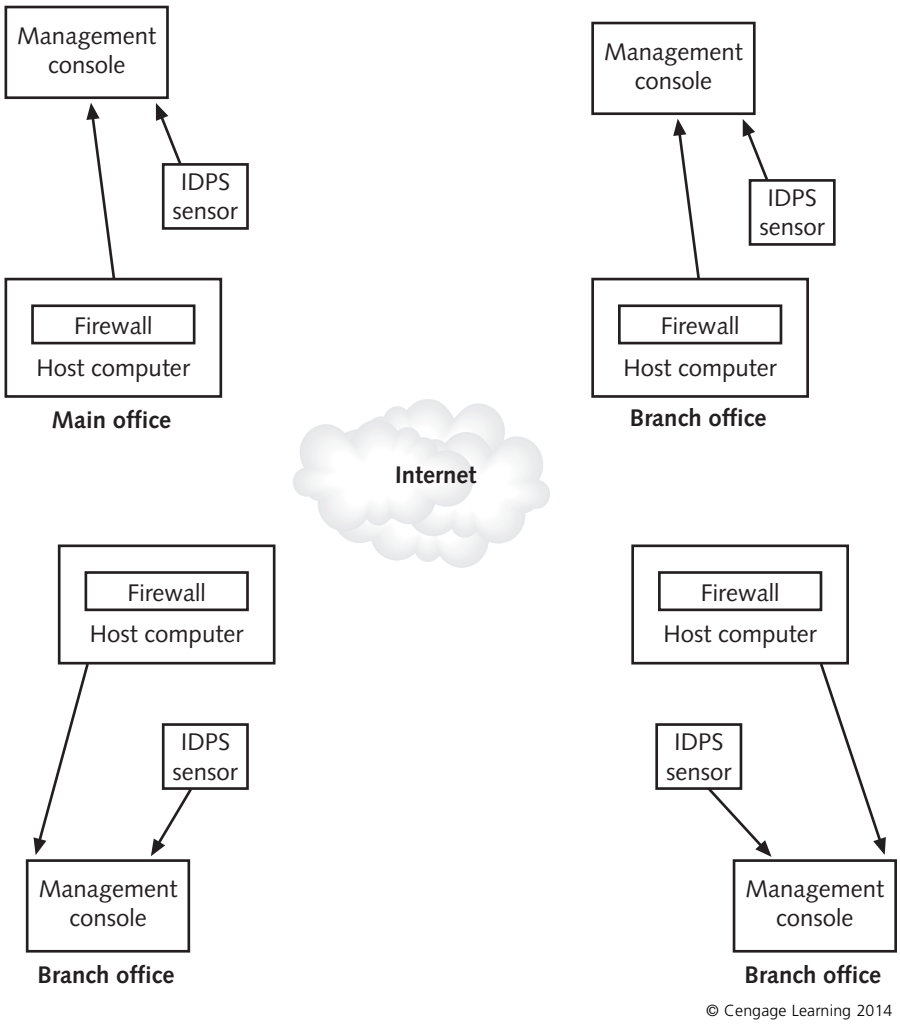

© Cengage Learning 2014

Figure 14-3 Distributed data collection

Evaluating IDPS Signatures

You do not monitor and evaluate network traffic to create a large quantity of data to analyze. Instead, you want to gather evidence that indicates whether your IDPS signatures are working well or need to be updated. A variety of IDPS vendors are available, each with

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

its own set of signatures for suspicious events. How do you evaluate signatures for each one? Various attempts have been made to standardize the evaluation and comparison of signature sets, but none have been adopted widely enough to be effective.

Therefore, the process of updating IDPS signatures varies depending on the vendor. Check the vendor's Web site for complete instructions and visit it often to download new signatures. Your organization might also hire a security management firm on a contract basis to handle these routine tasks.

Managing Change

The process of changing a procedure, a network defense component, or another administrative matter does not require a long, complex chain of events. However, you can encounter serious problems if you make an abrupt change that has an impact on the way employees work. The impact could even prove more harmful than the original problem. For example, suppose that you receive an alert from an IDPS and determine that you need to restrict access to a database server to selected administrative staff because an intruder nearly gained access to it. If you make the change without assessing its impact or without notifying all employees affected by the change, you are likely to be flooded with a stream of protests because employees will not be able to access files.

In an administrative sense, it is best to make sure that significant changes are carried out systematically and smoothly. Change management involves modifying equipment, systems, software, or procedures in a sequential, planned way. The process should include an assessment of the change's impact and a decision on whether to make the change based on that impact. Consider using change management for the following types of changes:

- Significant changes to firewall or IDPS rules that affect users
- New VPN gateways
- Changes to access control lists
- New password systems or procedures

The full process of evaluating and making a change is shown in Figure 14-4.

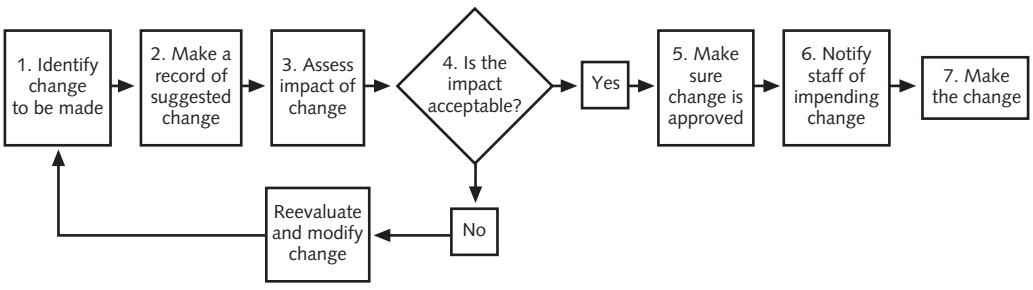

© Cengage Learning 2014

Figure 14-4 Significant changes should be evaluated and communicated systematically

The process might seem complex, but being systematic and thorough helps an organization run smoothly. You ensure that everyone is informed, you get the approval of managers, and

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

you make sure that the change should take place. The following list describes the process in Figure 14-4 in more detail:

1. Identify the change to be made. For example, you might decide to make a change after reviewing log files, after an intrusion, or after an alert message.
2. Create a written record of the suggested change. Describe the circumstances that prompted you to consider the change so you can communicate them to management and staff later.
3. Assess the impact of the change. Will the change inconvenience employees by making it more difficult to access resources they need? Will the change slow down network traffic by creating another layer of security? Think through all possible implications of the change.
4. Decide whether the impact of the change is acceptable. Discuss this question with the head of your department or your colleagues. You might also want to consult financial staff to determine whether the change has monetary implications. If the impact is not acceptable, reevaluate the change and determine whether a further modification will reduce the impact. If the impact is acceptable, proceed to the next step.
5. Make sure that the appropriate administrators approve the change. The approval process can add a considerable amount of time to the change. If you circulate a memo describing the change and a meeting is needed to discuss it, you might not be able to make the change for a week or more. If the change seems urgent and you think that a delay could create a security risk, try to talk to managers yourself and get approval immediately.
6. Notify staff of the impending change. After approval is given, tell all affected employees when the change will take effect. Give people several days or even a week to prepare for the change. However, if the change is needed urgently, make it immediately after receiving approval and tell people after the fact.
7. After notifying the affected employees and answering any questions, make the change.

You do not necessarily have to follow all the preceding steps before making every change. For example, to make a slight modification in a firewall or an IDPS rule, you do not need to tell the entire staff. The preceding procedure is needed only for significant changes that affect the way employees perform their daily work.

Strengthening Analysis: Security Auditing

Security auditing is the process of testing a network defense system. You can attempt break-ins to test the effectiveness of network defenses, for example. Recording and analyzing events such as logons, logouts, and file accesses can also supply useful information. You should also examine the security procedures in your organization, such as the ways it handles confidential information.

Groups known as **tiger teams** are assembled to actively test a network. Tiger team members usually have expertise in one or more areas of security, and each contributes to the overall effectiveness of the team. These penetration-testing teams can be used on a contract basis, but

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

you can also do testing yourself. You might have to accumulate data from several sources, such as the following:

- Packet filters
- Application logs
- Router logs
- Firewall logs
- Event monitors
- Host-based IDPS (HIDPS) logs
- Network-based IDPS (NIDPS) logs

One way to consolidate the data from these devices is to transfer (or "push") the information to a central database. Most IDPSs have a feature for transferring data; even the freeware IDPS program Snort allows you to transmit data to a database directly. At the very least, you should store the time of each incident, the data collected about each incident, the application affected, the operating system in use, the user ID, the process ID, and the log entry ID.

When using multiple security devices, you can accumulate so much data from log files that you need to manage it before it consumes your available storage space. Choose a time period for retaining detailed information from IDPS logs. Ninety days is a common choice. When the data is more than 90 days old, you can archive it to long-term storage, such as tape, DVD, or CD. Also consider paring data down to the most essential information before storing it.

Operational Auditing

In an **operational audit**, a company's IT staff examine system logs to confirm whether they are auditing needed information and not getting bogged down with unnecessary data. An operational audit should look for the following:

- Accounts that have weak passwords or no passwords
- Accounts assigned to employees who have left the company or user group
- New accounts that need to be checked against a list of authorized users

Financial institutions must undergo security audits regularly because of government regulations. These audits might include attempts at **social engineering**—attempting to trick employees into divulging passwords or other information. Auditors who work onsite in financial institutions even look in trash cans to see whether computer printouts of confidential information have been discarded without being shredded. Another auditing strategy is a **Tinkerbell program**, in which network connections are scanned and alerts are generated when connection attempts are made from suspicious sites or logons are attempted from a suspicious IDPS.

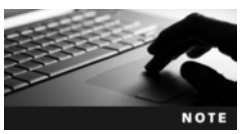

Tinkerbell is a monitoring program that scans incoming network calls and generates alerts when calls are received from particular sites or logins are attempted with a certain IDPS. The program was named after Project Tinkerbell, an experimental phone-tapping program that British Telecom developed in the early 1980s.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Because of its highly sensitive nature, information gathered in a security audit should be protected rigorously. (This information might include passwords, for instance.) If an unauthorized employee or intruder gained access to the files, network resources could be compromised. Your organization should have a clearly defined plan for handling and protecting audit data and should follow it accordingly.

Independent Auditing

In an **independent audit**, an outside firm inspects a company's audit logs to ensure that the organization is collecting the information it needs and not gathering unnecessary information that consumes system and network resources. The outside firm attempts to detect flaws or vulnerabilities in the entire system—not just in the IDPS but in other locations, such as files or applications. The audit firm might examine where security equipment is located, how well it is protected from unauthorized users or environmental disasters, and how thoroughly data is erased after disposal. Because the outside firm requires access to sensitive data, the auditors should sign a nondisclosure agreement (NDA) in which they promise not to release information to anyone outside the audit firm.

Auditing might uncover information about an organization's employees that they consider confidential. The company's legal staff should review the audit information and determine whether searching it constitutes an invasion of privacy. In government agencies, the concern over privacy increases: The Privacy Act of 1974 requires that citizens be notified when the government gathers information about them.

Strengthening Detection: Managing an IDPS

As your network grows, the amount of traffic and information on your IDPS and other security devices increases as well. To keep the IDPS running smoothly, you might need to modify the IDPS rule base and the amount of storage space available to the IDPS. This section briefly examines how to strengthen an IDPS by maintaining the current system and by changing or adding software or hardware.

Maintaining Your Current System

You do not necessarily need to add new systems or components to make an IDPS stronger. You can boost efficiency and strengthen detection by making your current resources work more efficiently. The following sections discuss how to maintain your system through backups and by managing accounts, IDPS rules, and users.

Backups You need to back up your firewall and IDPSs as a precaution. Store the backup in a secure location so you can restore systems if they become corrupted or intruders gain access to them. You should also keep backups of other security components:

- Routers
- Bastion hosts
- Servers
- Special-purpose devices

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Using automated backup software is advisable, especially for a large-scale network, because built-in backup software often does not have enough features to perform the tasks you need.

Managing Accounts Account management is another aspect of ongoing security maintenance. This often-neglected task involves adding new accounts, recovering old ones, and changing passwords. Make sure that user accounts are reviewed every few months. You must make sure that attackers have not added accounts and that terminated employees no longer have active accounts. Your security policy should include procedures that require the Human Resources Department to notify you whenever someone leaves the organization so you can disable or delete the account immediately.

Managing IDPS Rules Eliminate any unnecessary IDPS rules to reduce the amount of extraneous processing the IDPS must perform. The IDPS can then keep up with information that passes through fast networks, such as those that use gigabit-speed interfaces.

Keep your IDPS audit logs in a secure location so that attackers cannot tamper with them or erase evidence of their intrusions.

User Management A security user-awareness program consists of training procedures that ensure employees, contractors, and business partners understand the company's security policy and how it should affect their behavior. This program can improve security by teaching employees how to use the system more securely.

You can raise employee awareness in a variety of ways. For example, you can give a lecture on security within the organization and demonstrate how easily an intruder can crack weak passwords. You should also prepare booklets for employees to read and sign. By signing the booklets, employees indicate that they have read the security policy and understand its procedures and regulations.

Changing or Adding Software

Once or twice a year, IDPS software vendors usually release updated software. When a vendor releases an updated version of your command console software, find out what type of upgrade path is needed. Find out whether the new version requires working with new data formats and installing new supporting software.

Changing or Adding Hardware

Adding hardware to your system can be expensive, but the cost is probably minimal compared with the cost of lost data, eroded customer trust, and network downtime that can occur after an intrusion. You might be able to address your company's hardware needs just by adding a second network card to monitor the network, or you might need to consider more expensive options. You should consider adding consoles to reduce the **target-to-console ratio**—the number of target computers on your network managed by a single command console. You might also want to reevaluate the placement of sensors in your network. If you are not catching all the traffic on a network segment, consider moving a sensor to a new location or adding a sensor. You could also convert your host-based system into a hybrid configuration by adding a network sensor.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Strengthening Defense: Improving Defense-in-Depth

The principle of defense-in-depth (DiD) should guide you during maintenance of your perimeter security configuration. The DiD approach, as you learned in previous chapters, calls for security through a variety of defensive techniques that work together to block attacks. On the simplest level, DiD can be achieved with a personal firewall and antivirus software to protect a single computer. When applied to network services, DiD calls for maintenance of the following areas:

- *Availability*—Information is made available to authorized users as needed.
- *Integrity*—The information exchanged by network users is accurate.
- *Authentication*—Users prove their identity to other users or computers so they can exchange information.
- *Confidentiality*—Information can be read only by intended recipients.
- *Nonrepudiation*—Neither the recipient nor the sender of information can deny participating in a business transaction across the network.

The last item in this list—nonrepudiation—requires more explanation. Repudiation sometimes occurs in business transactions, particularly those that involve purchases or deliveries of supplies. Many transactions today are initiated electronically. For instance, company A might order 10,000 widgets from company B via an e-mail message. The order is placed, the goods are shipped, and the payment is made electronically. However, if company A receives the bill and then denies that it placed the order, it is said to have repudiated the transaction. Similarly, if company B denies having received the electronic payment, it repudiated the transaction. Nonrepudiation is the use of authentication to guarantee that both the sender and recipient participated in the message. Through public-key infrastructure, the sender and receiver are authenticated through the exchange of digital certificates. Both have electronic records that confirm the transaction actually took place, in addition to records of the date and time the transaction occurred.

Active Defense-in-Depth

Active defense-in-depth is a particularly strong implementation of the DiD concept. Instead of passively waiting for attacks to occur and then reacting, security personnel expect attacks to occur and try to anticipate them. Active DiD calls for multiple levels or layers of protection and security approaches that overlap one another. It also calls for defending against serious network threats first. The additional layers of protection address less serious threats.

Active DiD requires respondents to think creatively and counter all possible threats, whether they are familiar or unfamiliar. In an active approach, the method of defense changes based on the threat. Security personnel have the ability to "flex" based on the nature of the threat and where it occurs.

One layer of security is training. Security personnel should be training and learning continually to keep up with attacks and countermeasures, and to enable their organization to remain flexible when it comes to network defense. Organizations can take the following steps to create a training cycle for strengthening defenses:

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

1. *Training*—Security staff are trained in the network defense configuration.
2. *Perimeter defense*—After training, security staff can configure a perimeter defense for the network.
3. *Intrusion detection*—After defenses are arranged and put online, IDPSs alert security staff to potential vulnerabilities by sending alerts.
4. *Intrusion response*—Security staff evaluate the alerts and respond to block attacks.
5. *New security approaches*—Countermeasures are taken to reduce the number of false positives and true positives (actual attacks) that the IDPS encounters, which strengthens perimeter defense and helps security staff learn more about network vulnerabilities and defenses.

By using training as an ongoing process instead of an event that occurs only before hiring, an organization improves its overall security stance while employees improve their ability to prevent attacks.

Adding Security Layers

An active DiD posture accounts for the fact that a single network cannot be protected fully unless all interconnecting networks are protected. Your goal is to establish trust—to provide nonrepudiation for your network and ensure the integrity and confidentiality of information passing through it. To improve security, you might need to add new layers of security measures. For example, you might add an IDPS to supplement your firewall and VPN. If you already have an NIDPS, you might want to add an HIDPS or a hybrid IDPS to strengthen detection and response.

To see how this strategy works, consider how the U.S. Navy implemented it. In an article titled "Defense-in-depth: Security for Network-Centric Warfare," Captain Dan Galik described the layers or "zones" of defense created to protect one of the U.S. Navy's intranets:

- A firewall
- Encryption
- Virus protection to filter out potentially harmful e-mail attachments
- Authentication
- Intrusion detection
- Access control
- Information integrity through SSL and IPsec
- Auditing

In addition, four defensive zones were created to protect users and communication between zones. The following zones allowed protection to extend from the network perimeter to users' desktops:

- Zone 1 included users' desktops, which were protected through passwords, access control lists, VPN encryption, and virus protection.
- Zone 2 was described as a single "community of interest" or group of computers joined in a network. Routers, firewalls, virus protection software, and VPN encryption protected this zone.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Zone 3 encompassed multiple communities of interest so that communications between them were secure. Protection methods included network intrusion detection, firewalls, VPN encryption, and virus protection.
- Zone 4 was the outermost zone, which was the boundary between a Navy information system and a public network, such as the Internet. Defensive mechanisms at this level included DNS servers, VPN encryption, firewalls, and authentication for routers.

By dividing its communication needs into separate systems and relying on multiple security methods, the Navy achieved effective external security. However, this system does not always provide strong defense from vulnerabilities, such as accidental misuse of resources or malicious attacks from insiders. DiD can often be improved by placing more emphasis on access control, user education and awareness, and increased levels of auditing.

Strengthening Performance: Keeping Pace with Network Needs

IDPS performance refers to its ability to capture packets and process them according to the rule base. Ideally, an IDPS captures all the packets that reach it, sends alarms for all suspicious packets, takes action to protect the network, and allows legitimate packets to pass through to the internal network. This level of performance, however, can be hampered by several factors, as you learn in the following sections.

Managing Memory

IDPS performance depends largely on the number of signatures it has to review when receiving a packet. For signatures of events that consist of only a single packet, memory requirements are almost nonexistent. However, for signatures that require a sequence of packets to reach the IDPS, the IDPS needs to maintain the connection state in memory. Memory is also needed to store information in cache and for any databases that contain IDPS configuration settings. The primary consideration is that the IDPS should have more than the minimum amount of RAM to maintain state information, which allows it to thwart attacks that occur over extended periods.

Managing Bandwidth

Firewalls, packet sniffers, and IDPSs must be able to process data as quickly as it moves through the network. If your network has a gigabit Ethernet gateway but your IDPS runs on a host with only a 100-MB Ethernet connection, data passes through more slowly than it should. For example, an IDPS should be capable of handling 50 percent of bandwidth utilization without losing the capacity to detect problems. Intrusion detection often begins to break down if your bandwidth use exceeds 80 percent of network capacity.

You can run Performance Monitor in Windows or the Top utility in UNIX or Linux to gauge how much overhead your IDPS is consuming. Make sure host computers that provide

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

network sensors and IDPS management consoles run as fast as possible. If performance lags, upgrade immediately to a faster computer or network card.

Managing Storage

Some intrusions take place over long periods, so storing a sizeable amount of historical data from IDPS log files can be essential to tracking long-term attacks. A key aspect of managing an IDPS is ensuring that enough storage space is available for current IDPS data and making sure that data is archived or deleted securely when it is no longer needed. Typically, a gigabyte or more of space is needed to store IDPS data.

Clear out storage media when they are full and the information on them is no longer needed; this is an important and often overlooked element of network security. To make sure that intruders, criminals, or malicious employees cannot recover this data, you need to shred documents and files completely.

For systems that store highly sensitive information, you might want to physically remove and degauss the physical drives. **Degaussing** is the process of magnetically erasing an electronic device, such as a monitor or a disk, to remove any stray data or magnetic fields. Simply deleting or erasing files does not completely remove all information from the disk. When you use a delete or erase command, you simply remove pointers that tell the computer which clusters are storing the information. That way, the computer can reuse those sectors if needed. However, until the sectors are overwritten, the information in them can be recovered.

Staying Informed About Security Trends

To stay current with the latest trends in network security, you need to maintain a list of industry contacts, just as you maintain software and hardware configurations. Visit security-related Web sites and hold online discussions with other professionals in the field. Make sure management understands that you need to stay informed by participating in mailing lists and newsgroups, subscribing to trade publications, and taking certification tests to keep up with the fast-changing field of network security.

Web Sites

You will have no trouble finding Web sites that report the latest news on virus outbreaks and security breaches that affect prominent corporations. The challenge is to choose one or two sites you can scan daily or every few days for security issues that might affect your own network. You might visit sites such as the following:

- Center for Internet Security (*www.cisecurity.org*)
- SANS Institute (*www.sans.org*)
- CERT Coordination Center (*www.cert.org*)

Mailing Lists and Newsgroups

Mailing lists often provide more current information about security issues and vulnerabilities than Web sites or periodicals. These lists provide firsthand information from

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

security professionals about the problems they face, their opinions about security problems, and suggested solutions to those problems. Investigate the following mailing lists as a starting point:

- BugTraq (*www.securityfocus.com/archive/1*)
- Firewalls Mailing List (*www.securityfocus.com/archive/129/description*)
- SANS NewsBites (*www.sans.org/newsletters/newsbites/*)

Be careful about the messages you post on newsgroups or mailing lists. For example, do not provide specific details of security breaches your company has experienced; attackers might be monitoring the newsgroup or mailing list and find out that your network has been vulnerable. Do not provide specific details about your network configuration or identify your company. Consider using an e-mail address other than your company address for posting to newsgroups and mailing lists.

Certifications

Many certifications need to be renewed periodically through retesting. Make sure your organization's management understands that certifications benefit the company. Visit the following certification Web sites to stay current with tests you have taken and to learn about tests you might take in the future:

- Security Certified Program (*www.securitycertified.net*)
- International Information Systems Security Certification Consortium (*www.isc2.org*)
- CompTIA (*www.comptia.org*)
- GoCertify (*www.gocertify.com*)

Be aware of when you need to renew your current certifications. Ask your supervisor whether the company will help defray the cost of study materials and testing, and if you can have time off to study and travel as needed.

Chapter Summary

- Security event management involves accumulating data from a wide range of security devices by means of a coordinated program. This program includes monitoring alert and event logs produced by security devices and operating systems, and collecting data from multiple sensors through a centralized or distributed system. The program also requires reviewing IDPS attack signatures to make sure they are up to date.
- An often overlooked aspect of event management is the need to change a procedure in a systematic way. Change management describes the modification of systems or procedures using a disciplined approach that includes management approval and staff notification of an impending change before it occurs.
- Security auditing is used to test the effectiveness of network defenses after they have been established. In an operational audit, a company's IT staff examine system logs and look for vulnerabilities, such as weak passwords or unnecessary user accounts. An independent audit is performed by an outside auditing firm.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

- Another aspect of ongoing security maintenance is managing an IDPS to keep it running smoothly. You should make backups, manage user accounts, eliminate unnecessary IDPS rules, and add hardware or software as needed. Also institute an awareness program to make sure that employees, contractors, and business partners understand and observe the security policy.
- By strengthening a network's DiD configuration, you improve overall network defense, ensure the availability and integrity of information, and provide for nonrepudiation. Active DiD calls for trying to anticipate and thwart attack attempts before they occur. Active DiD also requires training or adding layers of security.
- Keep pace with your network's security needs by ensuring that your organization has enough memory for an IDPS to process long-term attacks, enough network speed to capture and process all packets, and enough storage space for log and alert files.
- Make sure that files are deleted completely by "shredding" them electronically—in other words, overwriting them so that all traces are removed from the media where they were stored.
- You must stay up to date with the latest security breaches and virus outbreaks. Visiting security-related Web sites, joining mailing lists, subscribing to newsletters, and posting on newsgroups are good ways to keep informed. In addition, keep your security certifications up to date to maintain your level of expertise.

Key Terms

active defense-in-depth A strong implementation of the DiD concept in which security personnel expect that attacks will occur and try to anticipate them instead of passively waiting for attacks and then reacting.

centralized data collection A system in which an organization's event and security data is funneled to a centralized management console in the main office.

degaussing The process of magnetically erasing an electronic device, such as a monitor or a disk, to remove any stray data or magnetic fields.

distributed data collection A system in which data from security devices goes to a management console on its own local network.

event monitoring Reviewing alert and event logs produced by security devices and operating systems, and testing the network periodically to identify weak points.

independent audit An audit in which an outside firm inspects audit logs to ensure that an organization is collecting the information it needs and not gathering unnecessary information that consumes system and network resources.

operational audit An audit by an organization's own staff that examines system and security logs to analyze information about intrusions and other unauthorized accesses.

security auditing The process of checking the effectiveness of a network defense system by testing the system, analyzing event logs, and observing procedures.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

security event management program A program that gathers and consolidates events from multiple sources so that the information can be analyzed to improve network security.

social engineering A technique of tricking employees into divulging passwords or other information.

target-to-console ratio The number of target computers on a network managed by a single command console.

tiger teams Special teams assembled to test a network.

Tinkerbell program A program in which network connections are scanned and alerts are generated when logons are attempted from a suspicious IDPS or connection attempts are made from suspicious sites.

Review Questions

1. Which of the following is a type of security audit? (Choose all that apply.)
   a. automated
   b. independent
   c. centralized
   d. operational

2. Why is it important to protect the confidentiality of information you gather through auditing? (Choose all that apply.)
   a. Employee privacy could be compromised.
   b. The information might become corrupted when you store it.
   c. Intruders could discover passwords.
   d. Viruses could infect it.

3. When should you follow the procedure for carrying out change shown in Figure 14-4?
   a. when many employees will be affected by the change
   b. when the change is needed urgently
   c. whenever a change needs to be made to security configurations
   d. when the change will have a substantial impact

4. What is an auditing program in which current connections are scanned and alerts are generated after suspicious logon attempts?
   a. social engineering
   b. port scan
   c. event monitoring
   d. Tinkerbell program

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

5. Which of the following employees has primary responsibilities that include maintaining and strengthening network defenses?
   a. security incident response team leader
   b. computer security manager
   c. chief information officer
   d. security auditor

6. What is a realistic goal of ongoing security management? (Choose all that apply.)
   a. blocking all suspicious packets
   b. tracing all attacks
   c. tracing as many intrusion attempts as possible
   d. continually strengthening and modifying defenses

7. Which of the following describes a goal of a security event management program? (Choose all that apply.)
   a. consolidating events from multiple sources
   b. responding to events as quickly as possible
   c. conducting forensics to trace and prosecute offenders
   d. managing IDPS signatures

8. How can you gather information on a variety of security events and respond to it quickly?
   a. Assemble a large response team.
   b. Use distributed data collection.
   c. Automate data collection and analysis.
   d. Outsource security management.

9. Which of the following is an advantage of centralized data collection? (Choose all that apply.)
   a. reduced traffic through network gateways
   b. reduced administrative costs
   c. reduced software and hardware costs
   d. only one person needed to review data

10. Why would you choose distributed data collection rather than centralized data collection?
    a. to reduce traffic through gateways
    b. to reduce the load on security managers
    c. to reduce overall costs
    d. to reduce hardware and software costs

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

11. Before installing new signatures for an IDPS, what do you need to do?
    a. Back up the IDPS.
    b. Stop the IDPS.
    c. Change passwords.
    d. Double-check to verify whether new signatures are necessary.
12. What can happen if you change a security configuration too abruptly and without proper authorization? (Choose all that apply.)
    a. Employees might ignore the change.
    b. The change might surprise other security managers.
    c. You might be flooded with protests from employees.
    d. You could face disciplinary action.
13. The change management process might apply when which of the following occurs? (Choose all that apply.)
    a. New password logon procedures are needed.
    b. You need to block access to DMZ servers.
    c. A new VPN gateway is installed.
    d. You need to change a fragmentation rule in a packet filter.
14. Security auditing involves which of the following? (Choose all that apply.)
    a. reviewing log files
    b. reviewing hardware and software costs
    c. testing defenses
    d. rotating firewall logs
15. What is nonrepudiation?
    a. the ability of a system to authenticate users
    b. the ability to rely on information gained through a security audit
    c. a legal defense used by employees whose privacy has allegedly been violated
    d. the ability to validate transactions through electronic documentation

Hands-On Projects

Hands-On Project 14-1: Assembling Security-Related Bookmarks for Web Pages

Time Required: 20 minutes

Objective: Create a list of Web sites to use as resources for staying updated on the latest developments in security.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Description: As part of your ongoing security management program, you need to consult dedicated Web sites and other resources to remain informed about network defense issues. One way to ensure that you check these resources regularly is to assemble a list of bookmarks that you can access quickly. You can then set up an e-mail reminder to check for updated software and other security-related news. For this activity, you need a computer with an Internet connection and Internet Explorer 9 installed.

1. Start your browser and go to ***www.securityfocus.com***.
2. Click the **Favorites** icon (see Figure 14-5), click the **Add to favorites** arrow, and click **Organize favorites**.

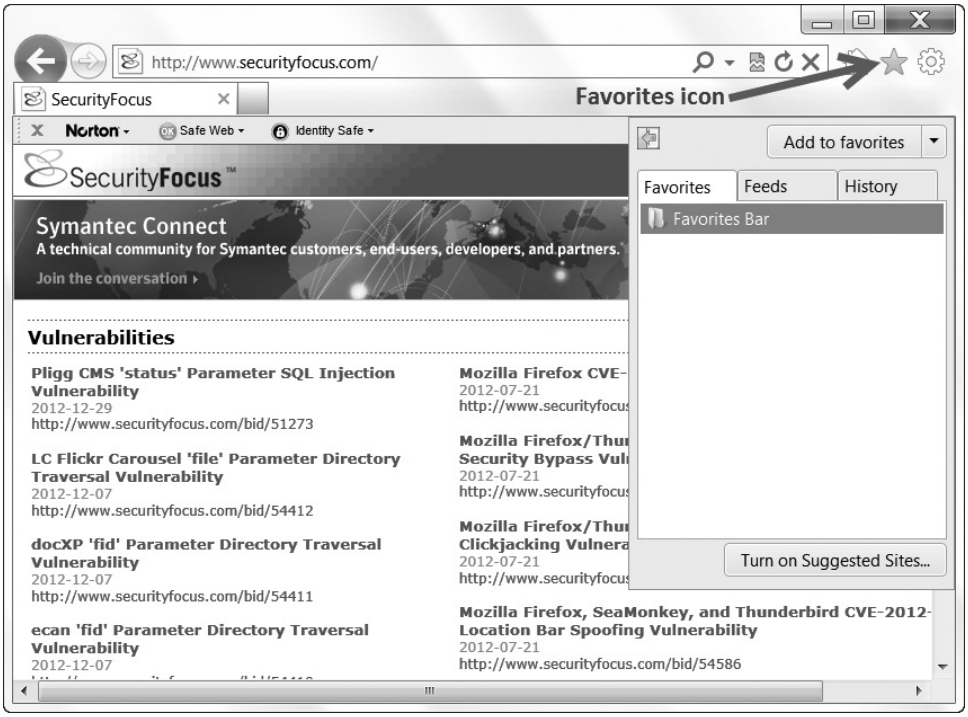

Source: Symantec Corp.

Figure 14-5 Internet Explorer favorites

3. Click **New Folder**, replace "New Folder" with the name **Security**, and then press **Enter**.
4. Click **New Folder**, replace "New Folder" with the name **News**, and then press **Enter**.
5. Click **Move**, click the **Security** folder, and then click **OK**.
6. Repeat Steps 4 and 5 to create two more new folders named **Conversations** and **Other Resources**.
7. Click **Close**.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

8. Add the Security Focus home page to the News folder. Click the **Favorites** icon, click **Add to favorites**, click the **Create in** arrow, click **News**, and click **Add**.
9. Near the top of the Security Focus home page, click **Join the conversation**. Add this page to the Conversations folder using the procedures in Step 8.
10. Scroll down the Security Focus home page to the Mailing Lists section. You could choose to subscribe to mailing lists, but in this step you add a mailing list to your favorites. Find the Penetration Testing heading and click **Complete Archive** below it.
11. Click the **Favorites** icon, click **Add to favorites**, and enter **Penetration Testing** in the Name text box over the default name. Verify that the Create in text box is set to Conversations, and click **Add**.
12. Return to the Security Focus home page, and repeat Steps 10 and 11 to add two more mailing list sites to the Conversations folder.
13. As an optional final step, add the URLs listed earlier in this chapter to the appropriate folders in your Security favorites folder. Specifically, add the URLs listed in the "Staying Informed About Security Trends" section.

Hands-On Project 14-2: Assembling Security-Related Bookmarks for RSS Feeds

Time Required: 20 minutes

Objective: Create a list of RSS feeds to use as resources for staying updated on the latest developments in security.

Description: Rich Site Summary (RSS) feeds allow publishers to syndicate blogs, news, and other types of content that might change frequently. You can obtain specific clients and RSS reader software or view feeds in Internet Explorer and other browsers. For this activity, you need a computer with an Internet connection and Internet Explorer 9 installed.

1. Start your browser and go to ***http://searchsecurity.techtarget.com/rss***.
2. On the SearchSecurity Web page under Security Wire Daily News, click the RSS icon shown in Figure 14-6.

Source: Mozilla Foundation

Figure 14-6 RSS feed icon

3. In the SearchSecurity: Security Wire Daily News window, click **Subscribe Now**. In the next window, click **Subscribe**.
4. You can now access this RSS feed from the Favorites icon in the Feeds tab.
5. Search the Internet for three more RSS feeds associated with topics in this book. Add the feeds to your feeds list.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Case Projects

Case Project 14-1: Information Security Certifications

To ensure that information technology workers maintain up-to-date knowledge about information security, many employers require periodic certification. An ever-growing list of information security certifications is offered by a variety of vendors. Some commonly sought certifications are administered by ISC^2, SANS, and CompTIA. U.S. Department of Defense directive DoD 8570 includes certification requirements for both DoD workers and support contractors. The directive defines employee position levels and certifications required at these levels. As of this writing, the latest version of the directive is 8570.01-M (*www.dtic.mil/whs/directives/corres/pdf/857001m.pdf*).

For this project, summarize the requirements of DoD 8570 in your own words. Include a definition of the three IAT (Information Assurance Technical) position levels. Describe the typical tasks performed by these workers, list the certification requirements for these levels, and include a summary of the objectives of each certification required at these levels.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

appendix A

Security Resources

Intrusion detection and network security are constantly changing fields. To stay current with the latest developments, you should consult the Web sites and other resources listed in this appendix. Many of the sites offer white papers, research papers, and other background information on topics such as firewalls, packet filtering, authentication, and encryption. You can find policy hints and resources, disaster-planning guides, and tools to help you do your job. You should also check these sites to learn about the latest threats. Information about bugs, security holes, and patches to plug them will be available online before you read about them in a book.

New threats surface daily, and you need to keep up with them. Table A-1 lists links to helpful Web sites that can enhance your skills and knowledge.

You should also join newsgroups and mailing lists for IT professionals as a way to network with your peers and learn from their experiences. Sometimes you might run into a problem that you cannot solve, despite poring through manuals and textbooks, searching the Internet, and running every test you can think of. Table A-2 lists places to start when looking for other professionals who might have a solution to your problem.

Web site	Description
Common Vulnerabilities and Exposures (CVE), *www.cve.mitre.org*	The CVE standard enables security devices to share information about attack signatures and other vulnerabilities so the devices can work together.
Symantec Security Response, *www.symantec.com/security_response*	Symantec maintains an extensive database of viruses. The site includes information about security incidents, but the emphasis is on protecting against viruses and other harmful code and eliminating them.
Internet Storm Center, *http://isc.sans.org*	This site, which is affiliated with the SANS Institute (also noted in this table), specializes in explaining how to respond to intrusions, incidents, and security alerts. A world map shows security breaches reported by geographic region. The site also includes a list of current attack trends, such as frequently attacked ports and recently reported malicious software.
DShield, *www.dshield.org*	This site is the home of the distributed intrusion-detection system, in which network administrators from around the world share firewall and intrusion-detection log information in an effort to track attack patterns.

© Cengage Learning 2014

Table A-1 Helpful sites for IT security professionals (*continues*)

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Web site	Description
The Center for Internet Security, *www.cisecurity.org*	This nonprofit organization is devoted to developing security standards it calls "benchmarks." Benchmarks are available for Linux, UNIX, and other operating systems.
SysAdmin, Audit, Network, Security (SANS) Institute, *www.sans.org*	This research and education organization focuses on network security. SANS conducts seminars and workshops on security.
The CERT Coordination Center, *www.cert.org*	This group, which is affiliated with the Carnegie-Mellon Institute, lists security alerts, incident notes, and vulnerabilities on its home page. CERT also offers tips and articles about network security and training courses.
Forum of Incident Response and Security Teams (FIRST), *www.first.org*	This group is a coalition of security incident response teams from government, commercial, and academic organizations that seek to promote rapid reaction to security incidents by coordinating communication and sharing information.
The National Institute of Standards and Technology (NIST), *www.nist.gov* NIST Computer Security Division, Computer Security Resource Center (CSRC), *http://csrc.nist.gov*	NIST is a nonregulated U.S. federal agency that develops and promotes measurements, standards, and technology to enhance productivity, facilitate trade, and improve the quality of life. The Information Technology Security and Networking Division addresses information security.
Internet Assigned Numbers Authority (IANA), *www.iana.org*	IANA assigns and maintains number assignments for the Internet, including port numbers and protocol numbers. IANA also coordinates DNS and IP addressing.
Internet Engineering Task Force (IETF), *www.ietf.org/home.html*	The IETF is an international body of network designers, vendors, operators, and researchers that works to advance the Internet's evolution and smooth operation.
IEEE Computer Society, *www.computer.org*	This organization is a major international membership association for computer professionals. A fee is charged for membership, but the organization offers online courses, information, and professional networking opportunities for members.
Network World Security Research Center, *www.networkworld.com/topics/security.html*	This long-established publication for networking insiders includes tabs for cybercrime, firewalls, unified threat management systems, IDPSs, and more.

© Cengage Learning 2014

Table A-1 Helpful sites for IT security professionals (*continued*)

Web site	Description
SecurityFocus Mailing Lists, *http://online.securityfocus.com/archive*	Security Focus maintains security-related mailing lists for topics ranging from intrusion detection to firewalls to honeypots. One of its best features is the ability to search archived messages by topic without having to subscribe. However, by joining a list, you can get news daily.
SANS Computer Security Newsletters and Digests, *www.sans.org/newsletters*	The set of newsletters published by the SANS Institute includes a weekly *NewsBites* publication and a weekly Security Alert Consensus that lists current security threats and countermeasures.

© Cengage Learning 2014

Table A-2 Information security news sources

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Security Certification Sites

The following sites offer certifications that can be invaluable for finding employment in network security:

Global Information Assurance Certification (GIAC)—The GIAC Web site (*www.giac.org*) provides information about the SANS Institute certification exams. Programs range from the entry-level Basic Information Security Officer certification to more specialized certifications, such as GIAC Certified Firewall Analyst (GCFW).

The International Information Systems Security Certification Consortium (ISC^2)—This international nonprofit organization (*www.isc2.org*) is dedicated to maintaining a common body of knowledge on security. ISC^2 prepares and administers two of the most common certifications in network security: Certified Information Systems Security Professional (CISSP) and Systems Security Certified Practitioner (SSCP).

CompTIA Certification—The Computing Technology Industry Association (*www.comptia.org/home.aspx*) is best known for the A+ PC technician series of certifications. The CompTIA Security+ Certification exam is available to establish fundamental security competency in firewalls, encryption, and intrusion detection.

The Security Certified Program—The Security Certified Program (*www.securitycertified.net*) is the vendor-neutral administrator of the Security Certified Network Professional (SCNP) and Security Certified Network Architect (SCNA) exams. The lower-level certification, SCNP, is comparable in content and skill level to the SSCP certification from ISC^2.

Certified Wireless Security Professional—This certification is administered by the Certified Wireless Network Professional organization, which provides vendor-neutral enterprise Wi-Fi certification and training. The certification requires that you pass two exams: one on wireless network administration and one on wireless security skills. The exams cover wireless security threats, assessment, design, architecture, and policy. For details, go to *www.cwnp.com/certifications/cwsp*.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Glossary

absorption A type of loss in which certain materials absorb the RF signal.

acceptable use policy A policy that defines acceptable and unacceptable uses of company resources.

access control entry (ACE) An individual rule in an ACL.

access control list (ACL) A group of rules that permit or deny passage of packets through a router based on parameters such as source and destination IP address, port address, protocol time to live, IP options, and TCP flags.

access point (AP) A wireless device that acts as the base station for a wireless network segment and serves as the link between wired and wireless segments.

accountability The ability to track an attempted attack or intrusion back to its source.

acknowledgement (ACK) A type of control frame that a receiving station sends when a packet is received successfully with no errors.

active attacks Attacks that attempt to gather information for subsequent attacks by using active scanning—sending probe request frames on each available channel.

active defense-in-depth A strong implementation of the DiD concept in which security personnel expect that attacks will occur and try to anticipate them instead of passively waiting for attacks and then reacting.

active scanning The process of a wireless station sending a probe request frame on each available channel and waiting for a probe response frame from available APs.

Address Resolution Protocol (ARP) A networking protocol that resolves IP addresses to MAC addresses.

Advanced Encryption Standard (AES) The current U.S. government standard for cryptographic protocols, AES uses the Rijndael algorithm with key sizes of 128, 192, or 256 bits and a fixed block size of 128 bits. AES is a symmetric block cipher used for encryption in IEEE 802.11i.

agent An IDPS component that monitors traffic on a specific host.

amplitude The maximum departure of a wave from its undisturbed state; the height of the carrier wave.

amplitude modulation (AM) An analog modulation method in which the height of the carrier wave is changed.

amplitude shift keying (ASK) A binary modulation technique similar to AM that changes the height of the carrier wave. With ASK, the 1 and 0 bits do not both have carrier signals; the 1 bit has positive voltage and the 0 bit has zero voltage.

analog RF signal A continuous RF wave that oscillates between positive and negative voltage.

anomaly detection A type of detection system that sends an alarm when it detects an event that deviates from behavior defined as normal in the IDPS profile.

antenna A device used to transmit or receive radio waves. One end of an antenna usually points into the air and the other end is fixed to the ground directly or via a grounded device. Antennas are generally made of copper wire or other conductive material.

antispoofing Measures to address digital attacks in which false source and destination addresses are used to bypass a router's packet filters.

anycast An address created automatically when a unicast address is assigned to more than one interface. Anycast addresses are assigned from unicast address ranges and have the same scopes as unicast addresses.

anycast addressing A network addressing scheme that allows DNS services to be decentralized among a group of servers, regardless of their location.

ARP table A network system tool that lists the MAC and IP address resolutions of other devices on the network, making the resolution process more efficient.

assets The hardware, software, and informational resources you need to protect by developing and implementing a comprehensive security policy.

association A two-step process of being accepted into a wireless network. First, a station listens for beacon frames to locate a network to join, and then it goes through the authentication process. Second, the station sends an association request frame; if it is accepted, the station receives an association ID and connection information. A station can be associated with only one network at a time, and it must be authenticated before being associated.

asymmetric algorithm A type of mathematical formula that generates a key pair; ciphertext generated by one key can only be decrypted by the other key.

auditing The process of reviewing records of network computer activity; these records identify who is connecting to a computer, the resources being requested, and whether access is granted or blocked.

authentication The process of verifying the identity of a user, computer, or service.

Authentication Header (AH) An IPsec protocol that provides authentication of TCP/IP packets to ensure data integrity.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

authentication, authorization, and accounting (AAA) logging The process of verifying users' identities, establishing their level of access, and recording their activity with the use of a RADIUS or TACACS+ server.

auxiliary (AUX) port The port on a Cisco device that permits dial-in access to the device.

availability The assurance that authorized users can access resources in a reliable and timely manner.

back doors Ways of gaining unauthorized access to a computer or other resource, such as an unused port or terminal service.

banner A message, usually a warning about appropriate use, presented to users of a digital system before authentication.

basic authentication An authentication method that uses a username/password pair to verify the identity of the user requesting access.

Basic Service Area (BSA) The geographical area of coverage for a WLAN.

Basic Service Set (BSS) An IEEE-defined WLAN configuration in which a group of wireless devices is served by a single AP.

bastion hosts Computers that are accessible to untrusted hosts.

biometrics A method of authenticating a user using physical information, such as retinal scans, fingerprints, or voiceprints.

block cipher A type of encryption algorithm that encrypts groups of cleartext characters.

botnets Networks of computers owned by unsuspecting victims of exploitation and controlled from a central system.

broadcast A communication sent to all hosts on a specific network.

buffered logging The process of outputting logged data to a Cisco router's RAM.

carrier signal The signal on which a radio transmission is sent.

centralized data collection A system in which an organization's event and security data is funneled to a centralized management console in the main office.

challenge/response authentication An authentication method in which one party presents a question, called the challenge, and the other party must provide the correct response, usually a password, to be granted access.

checksum A simple error-checking procedure used to determine whether a message has been tampered with in transit.

chipping code The bit pattern used in direct sequence spread spectrum. See also *direct sequence spread spectrum (DSSS)*.

ciphertext Unreadable text, programs that do not execute, and graphics you cannot view.

Cisco Internetwork Operating System (IOS) The operating system used on Cisco devices such as routers, switches, and firewalls.

Classless Interdomain Routing (CIDR) An IP address notation method that uses a slash (/) followed by the number of masked bits for an address—for example, 192.168.6.5/27 instead of an IP address of 192.168.6.5 and a subnet mask of 255.255.255.224.

cleanup rule A packet-filtering rule that comes last in a rule base and covers any packets that have not been covered by preceding rules.

clear to send (CTS) A type of control frame sent in response to a request to send (RTS) frame; it gives the sender clearance to begin transmitting packets.

client-to-site VPN A type of VPN connection that makes a network accessible to remote users who require dial-up access; also called a *remote access VPN*.

co-channel interference Interference from two or more wireless access points transmitting with the same frequency.

command console Software that provides a graphical interface to an IDPS.

Common Gateway Interface (CGI) scripts Scripts used to process data submitted over the Internet.

Common Vulnerabilities and Exposures (CVE) A standard that enables security devices to share information about attack signatures and other vulnerabilities so that they can work together to protect networks.

confidentiality Preventing intentional or unintentional disclosure of data during its creation, transmission, and storage.

console (CON) port The port on a Cisco device that permits direct physical access from a nearby computer using the serial RS-232 protocol.

control frame A type of MAC frame used to help deliver data frames between stations and to control access to the medium.

convergence A state in which all routers on a network have up-to-date routing tables.

cost-benefit analysis A technique for comparing the costs of an investment with the benefits it proposes to return.

crackers Hackers who break into systems with the intent of doing harm or destroying data.

cryptanalysis The study of breaking encryption methods. Some common attack methods against cryptographic systems include differential and integral cryptanalysis, random number generator attacks, side channel attacks, and XSL attacks.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

cryptographic primitives Modular mathematical functions that perform one task reliably. They form the basic building blocks of modern cryptography.

cryptographic protocol A detailed description that incorporates standardized requirements and guidelines for key generation and management, authentication, encryption, hashing functions, nonrepudiation methods, and other aspects of message security.

cryptography The process of converting plaintext into ciphertext by using an encoding function.

Data Encryption Standard (DES) An older protocol composed of a 16-round Feistel network with XOR functions, permutation functions, 6×4 S-box functions, and fixed key schedules. DES generates 64 bits of ciphertext from 64 bits of plaintext by using a 56-bit key.

data frame A type of MAC frame that contains the TCP/IP packet sent over a wireless network or between wireless devices. It carries data from higher-layer applications, such as printer control data or Web pages.

datagram A discrete chunk of information; each datagram contains source and destination addresses, control settings, and data.

dBm (decibel-milliwatt) A measure of relative signal strength; 1 mW = 0 dBm.

decibel The unit of measurement of the difference in strength between two signals.

default key A wireless station or AP's key for encrypting messages for transmission. The default key does not have to be the same on every station, but the same key must be used for both encryption and decryption.

default routes Manual entries in a router's routing table that specify how to handle packets that are not routable.

defense in depth (DiD) A strategy for achieving information security that uses multiple layers of defense.

degaussing The process of magnetically erasing an electronic device, such as a monitor or a disk, to remove any stray data or magnetic fields.

demilitarized zone (DMZ) A semitrusted subnet that lies outside the trusted internal network but is connected to the firewall to make services publicly available while still protecting the internal LAN.

diffraction A phenomenon that is similar to refraction, except that the signal bends around objects in its path; usually, these objects have a rough surface. See also *refraction*.

diffused IR transmission An infrared transmission that relies on reflected light rather than a line-of-sight connection between the IR emitter and detector.

digital RF signals RF signals that oscillate between positive and negative voltage but are broken into discrete segments with many starts and stops in the data stream. By contrast, analog transmissions have a constant stream of data.

digital signature A method of verifying nonrepudiation and integrity in messages.

direct sequence spread spectrum (DSSS) A wireless modulation technique that uses an expanded redundant code to transmit each data bit. The spreading code (key) is the chip sequence used to represent message bits, and the key is applied at the data level.

directed IR transmission A form of IR transmission that requires the emitter and the detector to be pointed directly at one another.

discretionary access control (DAC) An access control method that allows users to share information with other users; however, the risk of unauthorized disclosure is higher than with the MAC method.

distance-vector routing protocol A protocol that uses simple distance measurements such as hops to determine the best route for a packet.

distributed data collection A system in which data from security devices goes to a management console on its own local network.

distributed denial of service (DDoS) attack An attack in which many computers are hijacked and used to flood the target with so many false requests that the server cannot process them all, and normal traffic is blocked.

DNS Security Extensions (DNSSEC) A protocol designed to improve DNS security by using cryptography to ensure DNS integrity and authentication.

Domain Name System (DNS) A network service that provides a database for the resolution of fully qualified domain names to IP addresses, and of IP addresses to fully qualified domain names.

dual-homed host A computer configured with more than one network interface.

dynamic routes Entries in a router's routing table that are generated automatically by communication through protocols with other routers on the network.

electromagnetic (EM) radiation Electromagnetic energy traveling as a self-propagating wave and spreading out as it moves.

electronic assets Word-processing documents, spreadsheets, Web pages, and other documents on your network computers.

EM spectrum All forms of EM radiation collectively. From highest energy to lowest, EM radiation forms are gamma

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

rays, x-rays, ultraviolet light, infrared light, visible light, microwaves, and radio waves.

Encapsulating Security Payload (ESP) An IPsec protocol that encrypts the header and data components of TCP/IP packets.

encapsulation The process of enclosing a packet within another packet that has different IP source and destination information to ensure a high degree of protection.

encryption algorithm A precise set of instructions that provides an encoding function for a cryptographic system or generates output for use in additional operations.

Equivalent Isotropically Radiated Power (EIRP) The power radiated by a wireless system's antenna, including the intended transmission power and the unintended gain.

escalated Increasing an intrusion response to a higher level.

escalation procedure A set of roles, responsibilities, and measures taken in response to a security incident.

event horizon The entire length of an attack, from the first packet the IDPS receives to the last packet needed to complete the attack.

event monitoring Reviewing alert and event logs produced by security devices and operating systems, and testing the network periodically to identify weak points.

exclusive OR (XOR) function A cryptographic primitive based on binary bit logic and used as a linear mixing function, combining values for use in further computations.

exposure Vulnerability to loss resulting from the occurrence of a threat, such as disclosure, destruction, or modification of information resources. Exposure increases with the presence of multiple threats.

Extended Service Set (ESS) An IEEE-defined WLAN configuration in which multiple BSS networks are connected through a common distribution system.

Extensible Authentication Protocol (EAP) A group of extensible management protocols used in IEEE 802.1x. EAP includes a method of secure key exchange, and wireless stations use it to request port access.

extranet A private network that a company sets up as an extension of its corporate intranet so contractors, suppliers, and external partners can access a limited portion of the network infrastructure.

fading A form of interference/power loss in which the signal is reflected or refracted by the atmosphere, ground, or other obstacles, and the reflected path interferes with the direct path.

failover firewall A backup firewall that is configured to switch on if the current firewall fails.

false negatives Attacks that are not detected by an IDPS.

false positives Alarms generated by legitimate network traffic rather than actual attacks.

fast basic service set transition (FT) A method specified in 802.11r that allows decreased latency in the authentication of a wireless client with an AP to which it is roaming.

fault tolerance The ability of an object or a system to continue operating despite a failure.

Federal Communications Commission (FCC) The primary regulatory agency in the United States for wireless communication.

Feistel network A cryptographic primitive that forms the basis of many symmetric algorithms. Feistel networks combine multiple rounds of repeated operations, such as processing cleartext input with XOR functions. A key schedule is used to produce different keys for each round.

firewall Hardware or software configured to block unauthorized access to a network.

firewall appliances Hardware devices with firewall functionality.

firewall policy An addition to a security policy that describes how firewalls should handle application traffic, such as Web or e-mail applications.

footer A section sometimes added to a TCP/IP packet to indicate the end of the packet.

Fourth Amendment An amendment to the U.S. Constitution that provides protection from illegal search and seizure.

fragmentation The division of packets into smaller sizes to accommodate routers with frame size limitations.

free space path loss The loss of signal strength resulting from the dispersion of the signal over distance.

frequency The number of times an event occurs in a specified time period; the rate at which an electrical current alternates, creating different radio transmissions.

frequency division multiplexing (FDM) An RF signaling method that enables multiple users to share a communication channel. Base signals are modulated on different carrier waves and combined to form a composite signal.

frequency hopping spread spectrum (FHSS) A modulation technique in which the range of frequencies used for transmission changes during the transmission. The time spent on a particular frequency is called the dwell time.

frequency modulation (FM) An analog modulation technique that changes the number of waves used to represent one cycle, so that the number of waves representing a 1 bit is greater than the number of waves representing a 0 bit.

frequency shift keying (FSK) A binary modulation technique that changes the frequency of the carrier signal.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Fresnel zone An area in which you can determine RF signal strength, identify potential obstacles, and determine the potential for multipath distortion between sending and receiving antennas.

fully qualified domain names (FQDNs) Complete DNS names of computers that include the computer name, domain name, and domain name extension, such as *www.cengage.com*.

gain The positive difference in amplitude between signals.

Generic Routing Encapsulation (GRE) A nonproprietary tunneling protocol that can encapsulate a variety of Network layer protocols.

hacker Anyone who attempts to gain access to unauthorized resources on a network, usually by finding a way to circumvent passwords, firewalls, or other protective measures.

hactivists Computer attackers with political goals.

hashing algorithms Sets of instructions applied to variable-length input (the message) that generate a fixed-length message digest representing the input. Hashing algorithms do not provide confidentiality because they do not encrypt the message contents, but they do provide verification that a message has not been altered.

hashing functions Processes that a computer runs to verify message integrity by generating a hash value (also known as a message digest), which is a fixed-size string representing the original input's contents. Hashing functions are also used for error detection.

header The part of a packet that contains source and destination information and general information about the packet.

hopping code The sequence of changing frequencies in FHSS that determines the sequence and speed of frequency hops.

host identifier The part of an IP address that is unique to a computer in its subnet.

host-based IDPS (HIDPS) An IDPS deployed on hosts inside the network perimeter.

hybrid configuration A VPN configuration that combines characteristics of the mesh and star configurations.

hybrid IDPS An IDPS that combines the capabilities of an HIDPS and NIDPS for more flexibility and security.

ICMP ID number A number that identifies the ICMP packet so that the originating computer can make sure the response came from its original request.

IDPS management server An IDPS component that serves as the central repository for sensor and agent data.

IEEE 802.1x A standard developed to provide port-based access control on Ethernet LANs; it has been revised for wireless networks to incorporate authentication and key management.

IEEE 802.11 authentication The process by which an AP accepts or rejects a wireless device's connection attempt.

IEEE 802.11i A wireless security standard intended to replace the IEEE 802.11 WEP-based standard. It uses 802.1x authentication and AES for encryption. See also *Advanced Encryption Standard (AES)*.

IEEE Std. 802.11-2007 A series of wireless networking standards that specify transmission speeds, ranges, and other aspects of wireless networks.

independent audit An audit in which an outside firm inspects audit logs to ensure that an organization is collecting the information it needs and not gathering unnecessary information that consumes system and network resources.

Independent Basic Service Set (IBSS) A wireless network that does not use an AP; also known as ad-hoc mode.

infrared (IR) radiation A type of radiation with a shorter wavelength than radio waves but a longer wavelength than visible light.

infrared transmissions Transmissions that use infrared light pulses and that require an emitter and a detector. The emitter is usually a laser diode or LED. The detector is sometimes combined with the emitter and does just what its name implies: It detects IR transmissions and generates the corresponding electrical current for wired communication.

initialization vector (IV) In WEP's shared key encryption, a 24-bit value added to the beginning of a key to initialize the key stream generated by the RC4 algorithm.

inline sensor An NIDPS sensor positioned so that all traffic on the network segment is examined as it passes through. See also *network-based IDPS (NIDPS)*.

integrity The accuracy and consistency of information during its creation, transmission, and storage.

intentional radiator A device designed to generate radio signals, not including those from the antenna.

Internet Control Message Protocol (ICMP) A protocol that reports network communication errors to support IP communications. The ping command is a common troubleshooting utility based on ICMP.

Internet Key Exchange (IKE) A form of key exchange used to encrypt and decrypt data as it passes through a VPN tunnel. IKE uses tunnel method encryption to encrypt and then encapsulate packets for extra security. See also *tunnel method encryption*.

Internet Protocol Security (IPsec) A set of standard procedures that the Internet Engineering Task Force

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

(IETF) developed for enabling secure communication on the Internet.

Internet Protocol version 4 (IPv4) The IP addressing system currently in widespread use on the Internet, in which addresses are created with 32 bits (4 bytes) of data.

Internet Protocol version 6 (IPv6) A new version of IP that is gaining support among software and hardware manufacturers, and that will eventually replace IPv4; this version calls for 128-bit IP addresses.

Internet Security Association Key Management Protocol (ISAKMP) An IPsec-related protocol that enables two computers to agree on security settings and establish a Security Association so that they can use Internet Key Exchange. See also *Internet Key Exchange (IKE)* and *Security Association (SA)*.

Internet use policy A policy that defines how users can access and use the Internet and that specifies what rules apply to e-mail and other communications, such as instant messaging.

intrusion An attempt to gain unauthorized access to network resources and to compromise the integrity and confidentiality of network data or users' privacy.

intrusion detection and prevention system (IDPS) A network security measure that can consist of applications and hardware devices deployed on the network, hosts, or both to prevent, detect, and respond to traffic interpreted as an intrusion.

inverse mask The number in an access control list that specifies which part of an IP address is considered significant.

IP spoofing The process of inserting a false address into an IP header to make the packet more difficult to trace back to its source.

IPsec driver Software that handles the tasks of encrypting, authenticating, decrypting, and checking packets in an IPsec connection.

Kerberos An IETF standard for secure authentication of requests for resource access.

Key Distribution Center (KDC) A Kerberos component that holds secret keys for users, applications, services, or resources; KDC creates and distributes session keys by using symmetric cryptography.

key management A way to prevent keys from being discovered and used to decipher encrypted messages. One form of key management is to change keys frequently.

Layer 2 Tunneling Protocol (L2TP) A tunneling protocol derived from two older protocols (Cisco's L2F and Microsoft's PPTP L2TP) that encapsulates PPP packets and is usually combined with IPsec for improved security.

link-state routing protocol A routing protocol that uses complex calculations to determine the best route for a packet.

load-balancing software Software that prioritizes and schedules requests and then distributes them to servers in a server cluster based on each server's current load and processing power.

logging The process of recording activity on a computing device.

logic bomb Malware designed to be used at a specific time in the future or when a specified condition exists.

loss The negative difference in amplitude between signals.

macro A type of script that automates repetitive tasks in Microsoft Word or similar applications.

management frame A type of MAC frame used to establish and maintain communications between wireless devices, or between a wireless device and an access point.

mandatory access control (MAC) An access control method that defines an uncompromising manner for how information can be accessed. With the MAC method, all access capabilities are defined in advance.

many-to-one NAT A process that uses the source and destination TCP and UDP port addresses to map traffic between internal and external hosts. Many-to-one NAT is also called *Port Address Translation*.

maximum transmission unit (MTU) The maximum packet size that can be transmitted over a type of computer network. For example, the MTU of an Ethernet network is 1500 bytes.

mesh configuration A VPN configuration in which all participants in the VPN are connected to one another. This configuration is commonly arranged as a full-mesh or partial-mesh setup.

Message Authentication Code (MAC) A hashing algorithm that uses a shared secret key to generate a MAC tag for a message.

Message Digest 5 (MD5) A widely used hashing algorithm that produces a 128-bit hash value displayed as a 32-character hexadecimal number.

Message Integrity Check (MIC) A mathematical function used in WPA that replaces the cyclic redundancy check (CRC); MIC is designed to detect tampering in packets. See also *Wi-Fi Protected Access (WPA)*.

metrics Cost values that help routers assess the desirability of a link.

modulation The process of varying the wave form to transmit a signal, such as modulating between on and off to create 1s and 0s.

Monte Carlo simulation An analytical method that simulates a real-life system by randomly generating values for variables.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

multicast A transmission used for one-to-many communication, in which a single host can send packets to a group of recipients.

Multicast Listener Discovery (MLD) A core IPv6 protocol that enables IPv6 routers to discover multicast listeners on a directly connected link and to decide which multicast addresses are of interest to those nodes.

multipath A signal that has more than one path from the transmitter to the receiver.

multiple entry point configuration A type of VPN configuration in which multiple gateways are used, each with a VPN tunnel connecting a different location.

multiple-input multiple-output (MIMO) The use of multiple antennas operating at 2.4 GHz and 5 GHz, which allows for increased bandwidth in the 802.11n standard.

multiple-packet attacks Attacks that require a series of packets to be transmitted in order to execute.

named ACLs An advancement over extended and standard ACLs that allow filtering packets based on IP options and TCP flags, and that permit insertion and modification of individual ACEs.

narrowband An RF transmission on one frequency or a very small frequency range.

Neighbor Discovery (ND) A core IPv6 protocol used to resolve addresses, locate neighboring routers, and redirect hosts to better routes to reach destination addresses. ND uses ICMPv6 messages to manage node-to-node communications.

network access points (NAPs) Highly secure public facilities where commercial Internet backbones and ISPs exchange routing and traffic data.

Network Address Translation (NAT) A process by which internal hosts are assigned private IP addresses and communicate with the Internet using the public address of the NAT device.

network assets Routers, cables, bastion hosts, servers, and firewall components that enable employees to communicate with one another and other computers on the Internet.

network identifier The part of an IP address that a computer has in common with other computers in its subnet.

network security policy A policy that defines and establishes responsibilities for protecting the network and the information processed, stored, and transmitted on the network.

network-based IDPS (NIDPS) A set of IDPS components specialized for network use. An NIDPS examines traffic on network segments by using well-positioned sensors, management servers, a command console, and databases of signatures.

nonrepudiation The capability to prevent one participant in an electronic transaction from denying that it performed an action.

null packets TCP packets with no flags set.

one-to-one NAT The process of mapping one internal IP address to one external IP address.

open system authentication An 802.11 authentication method that relies on a station having an SSID that matches the network it is attempting to join; a matching SSID is the only criterion for acceptance. The security drawback of open system authentication is that SSIDs are transmitted in cleartext in management frames.

operational audit An audit by an organization's own staff that examines system and security logs to analyze information about intrusions and other unauthorized accesses.

orthogonal frequency division multiplexing (OFDM) A modulation technique that breaks a high-speed signal into smaller pieces and sends it in parallel across many lower-speed channels.

packet filters Hardware or software tools that allow or deny packets based on specified criteria, such as port, IP address, or protocol.

packet injection The insertion of crafted packets into network traffic.

packet monkeys A derogatory term for unskilled crackers or hackers who steal program code and use it in denial of service attacks instead of creating the programs themselves.

packet sniffer Software or hardware that monitors traffic into or out of network devices and captures information about each TCP/IP packet it detects.

pairwise keys Any keys used between a pair of devices in TKIP. See also *Temporal Key Integrity Protocol (TKIP)*.

pairwise master key (PMK) The TKIP key used to generate data encryption keys, data integrity keys, and session group keys, among others. This key is used only once at the start of a session.

pairwise transient key (PTK) The first TKIP key generated by the pairwise master key (PMK) and used for further key generation, according to the cryptographic algorithm in use.

passive attack An attack that uses passive scanning to gather information for later use in other attacks; also called *network reconnaissance*.

passive scanning The process of a wireless station listening on each available channel for an AP's beacon frame, and then moving on to the next channel without sending anything.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

passive sensor An NIDPS sensor that examines copies of traffic on the network; no actual traffic passes through the sensor. See also *network-based IDPS (NIDPS)*.

penetration testing The process of using a variety of tools and techniques to attempt to break into a network. Penetration testing, also called *ethical hacking*, is used legitimately as part of security audits to identify security vulnerabilities that attackers could exploit.

permissive policy A general approach to security that calls for a firewall and associated components to allow all traffic by default, blocking only specified traffic on a case-by-case basis.

permutation functions Bit-shuffling cryptographic primitives that reorder sets of objects randomly.

pharming A variation of phishing that intercepts traffic to a legitimate Web site and redirects it to a phony replica site.

phase modulation (PM) An analog modulation technique that changes the cycle's relative starting point when the bit being transmitted changes from 1 to 0.

phase shift keying (PSK) A binary modulation technique that starts and stops the signal to represent a binary digit.

phishing Using social engineering techniques via e-mail to trick users into providing personal information at the attacker's Web site. This site is designed to look like a legitimate business site.

physical security Measures taken to physically protect a computer or other network device from theft, fire, or environmental disaster.

ping sweep The act of sending a series of ICMP echo request packets in a range of IP addresses to see whether any computers respond.

plaintext Readable text, programs that execute, and graphics you can view.

point of presence (POP) ISPs ISP facilities that provide connectivity to the Internet for business, education, and home users.

Point-to-Point Tunneling Protocol (PPTP) A tunneling protocol used for dial-up access to a remote server.

polarization The plane on which radio waves propagate, or the orientation of the radio waves as they leave the antenna; usually the orientation is vertical or horizontal.

port An area in random access memory (RAM) reserved for the use of a program that "listens" for requests for the service it provides.

port scan An attempt to connect to a computer's ports to see if any are active and listening.

Power over Ethernet (PoE) A technology that sends power for the AP unit over unused wire pairs in standard twisted-pair Ethernet cabling.

power-save poll (PS-Poll) A type of control frame that a station sends after awakening from power-save mode; the frame indicates to the access point that it is ready to receive any frames the AP has buffered for it.

private key exchange In symmetric cryptography, the same key is used to encrypt and decrypt a message. Public-key Infrastructure is often used for private key exchange.

privileged access policy A policy that details additional access options and responsibilities of users with privileged access to resources.

probability The possibility that a threat will actually occur, as influenced by geographic, physical, habitual, or other factors that increase or decrease the likelihood of occurrence.

profiles Sets of characteristics that describe network services and resources a user or group normally accesses.

proxy server Software that forwards packets to and from the network being protected and caches Web pages to speed up network performance.

pseudorandom number generators (PRNGs) Cryptographic primitives used to generate sequences of numbers that approximate random values.

public key cryptography standards (PKCSs) A set of standards that RSA developed to provide standardization guidelines for cryptography. Many of these 15 standards have moved into the IETF standards track.

public key exchange In asymmetric cryptography, two keys are required: the public key and the private key. The public key used to encrypt the message is shared freely. The private key used to decrypt the message is kept secret.

radio frequency (RF) monitor mode A wireless NIC mode, equivalent to promiscuous mode in wired NICs, that allows a WNIC to capture packets without authenticating or associating with an AP or ad-hoc wireless network.

radio frequency (RF) waves A type of EM radiation used most commonly for wireless networks.

reflection A type of loss in which the RF wave bounces off certain materials.

refraction A phenomenon in which signals traveling through different air masses can be bent rather than traveling in a straight line. The signal changes in response to atmospheric conditions.

remote access and wireless connection policy A policy that defines necessary security measures before a remote desktop or wireless connection is added to an organization's network.

Remote Authentication Dial-In User Service (RADIUS) An authentication, authorization, and accounting service protocol.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Remote Procedure Calls (RPC) A standard set of communication rules that allows one computer to request a service from another computer on a network.

remote-access Trojans (RATs) Malicious programs that allow a remote attacker to take control of the target system.

request to send (RTS) A type of control frame that a station sends when it wants to transmit. The RTS frame sender requests that the medium be reserved long enough for the transmission to be completed and for an acknowledgement control frame to be received from the destination.

residual risk The risk remaining after countermeasures and defenses are implemented.

restrictive policy A general approach to security that calls for a firewall and associated components to deny all traffic by default, allowing only specified traffic on a case-by-case basis.

reverse firewall A device that filters outgoing connections.

Rijndael The encryption algorithm used in AES; a symmetric block cipher composed of 10 to 14 rounds of S-box and XOR functions. It supports 128-bit, 192-bit, or 256-bit keys and block sizes. Rijndael applies 10 rounds for 128-bit keys, 12 rounds for 192-bit keys, and 14 rounds for 256-bit keys.

risk The possibility of incurring damage or loss.

risk analysis A process of analyzing the threats an organization faces, determining which resources are at risk, and determining the priority of each asset.

risk management The process of identifying, choosing, and setting up countermeasures justified by the risks you identify.

roaming A wireless device moving between access points, or cells, without losing its ability to transmit and receive. Cells are interconnected so that as the wireless device moves between cells, it negotiates a new connection depending on signal strength.

rogue devices Wireless devices installed on a network without authorization or verified configurations.

role-based access control (RBAC) An access control method that establishes organizational roles to control access to information. The RBAC method limits access by job function or job responsibility.

role-based authentication A method of authentication that grants users limited system access based on their assigned role in the company and that defines the resources the role is allowed to use.

rollover cable A cable used to connect to a Cisco CON port from a computer's COM port via the serial RS-232 protocol.

route summarization The process of shortening the network identifier in a subnet mask to allow one network route to represent multiple network routes. (Also called *supernetting*.)

routed protocol A networking protocol used to transport data packets.

routing The process of transporting packets of information across a network from the source node to the destination node.

routing protocol A protocol used by routers to share their routing tables with each other.

routing table The information about accessible networks that is maintained by a router.

rule base The collection of rules that filter traffic at an interface of a firewall.

scan throttling The tactic of extending the time between scan probes to make the process more difficult to detect.

scanner A device that scans a network for open ports or other potential vulnerabilities.

scattering A type of loss in which small objects, such as raindrops, dust, or foliage, scatter the signal.

scopes Unicast addresses used in IPv6 to identify the application suitable for the address; scopes include global unicast, site-local unicast, and link-local unicast.

screened host A dual-homed host in which one interface is connected to an internal network and the other interface is connected to a router to an untrusted network.

screening router A router placed between an untrusted network and an internal network.

script kiddie A young, inexperienced computer programmer who spreads viruses and other malicious scripts and exploits weaknesses in computer systems using tools and techniques created by others.

scripts Executable code attached to e-mail messages or downloaded files that is used to infiltrate a system.

search warrant A legal document issued by a court that allows authorities to search a particular place for specific evidence. The warrant must detail what the search is seeking and where authorities are permitted to look for it.

Secure Hash Algorithm (SHA) A hashing algorithm that the NSA designed as a replacement for MD5. SHA-1 produces a 160-bit message digest.

Secure Shell (SSH) A protocol used to secure remote sessions by using asymmetric encryption.

Secure Sockets Layer (SSL) A protocol developed by Netscape Communications Corporation as a way of enabling Web servers and browsers to exchange encrypted information.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Security Association (SA) A designation for users, computers, or gateways that can participate in a VPN and encrypt and decrypt data by using keys.

security auditing The process of checking the effectiveness of a network defense system by testing the system, analyzing event logs, and observing procedures.

security event management program A program that gathers and consolidates events from multiple sources so that the information can be analyzed to improve network security.

security incident response team (SIRT) A group of people designated to take countermeasures when an incident is reported.

security policy A statement that defines the defenses configured to block unauthorized access to a network, acceptable use of network resources, an organization's response to attacks, and how employees should handle the organization's resources to prevent data loss or file damage.

security workstation A computer dedicated to providing firewall policies.

security-aware resolver A computing system that is compliant with DNSSEC and that attempts to resolve a fully qualified domain name to an IP address (or vice versa).

selective acknowledgements Acknowledgements that selected packets in a sequence have been received, as opposed to acknowledging every packet.

sensor An IDPS component that monitors traffic on a network segment.

server farm A group of servers connected in a subnet that work together to receive requests.

service set identifier (SSID) A unique identifier of up to 32 characters that serves as the wireless network name.

shared key authentication An 802.11 authentication method that uses a standard challenge-response process with shared key encryption. It relies on a wireless station having the correct key to encrypt the AP's challenge text as the criterion for acceptance. The method's main security drawback is that it encrypts the shared key using WEP, which is easily cracked.

side channel attack An attack method that exploits vulnerabilities of underlying hardware systems that leak data instead of exploiting vulnerabilities in a cryptographic algorithm.

signature A set of characteristics—such as IP numbers and options, TCP flags, and port numbers—used to define a type of network activity.

signature analysis The practice of analyzing and understanding TCP/IP communications to determine whether traffic is legitimate or suspicious.

signature detection The process of monitoring network traffic to detect unauthorized access attempts and sending notifications so that countermeasures can be taken.

signature files Files used by antivirus programs that contain patterns of known viruses and malware.

signatures Signs of possible attacks that include an IP address, a port number, and the frequency of access attempts; an IDPS uses signatures to detect possible attacks.

Simple Network Management Protocol (SNMP) A protocol used to monitor the status of network devices.

single entry point configuration A VPN configuration in which all traffic to and from the network passes through a single gateway, such as a router or firewall.

single-packet attack An attack that can be completed by sending a single network packet from client to host.

site survey An in-depth examination of a proposed wireless network site that is designed to determine AP placement and identify stray signals, noise, and obstacles.

site-to-site VPN A VPN that uses hardware devices, such as routers, to connect two networks; also called a *gateway-to-gateway VPN*.

sniffing Capturing network traffic during transmission.

SNMP trap logging A specific parameter of network activity to monitor.

social engineering A technique of tricking employees into divulging passwords or other information.

socket The end point of a computer-to-computer connection defined by an IP address and port address.

spear phishing A variation of phishing directed at specific users instead of using spam e-mail. Spear phishing is often directed at employees of a particular organization, for example.

split brain DNS architecture A network architecture that uses a single DNS domain with a DNS server on the organization's DMZ for Internet services and a DNS server on the internal network for service to internal hosts. All records that refer to internal hosts must be removed from the DMZ DNS server.

split DNS architecture A network architecture that divides DNS services between two servers: a public DNS domain with a server on the organization's DMZ for Internet services and an internal DNS domain with a server on the internal network for service to internal hosts.

split tunneling The term used to describe multiple paths. One path goes to the VPN server and is secured, but an unauthorized and unsecured path permits the user to connect to the Internet or some other network while still connected to the corporate VPN.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

spread spectrum An alternative to narrowband transmissions that spreads a narrow signal over a broader portion of the radio band.

standard ACLs Access control lists that are limited in filtering parameters to a source IP address or network.

star configuration A VPN configuration in which a single gateway is the "hub" and other participating networks are considered "rim" networks.

state information Data about a network connection that is typically kept in a state table.

state table A file maintained by stateful packet filters that contains a record of all current connections.

stateful autoconfiguration In IPv6, the method of dynamic host configuration that uses DHCPv6. The method is called "stateful" because the DHCPv6 client and server must keep their information updated to prevent addressing conflicts.

stateful packet filters Filters that are similar to stateless packet filters, except that they also determine whether to allow or block packets based on information about current connections.

stateful protocol analysis The process of maintaining a table of current connections so that abnormal traffic can be identified based on previous packets associated with the same transmission.

stateless autoconfiguration A feature of IPv6 in which a computer can connect to a network by determining its own IP address based on the addressing of neighboring nodes.

stateless packet filters Simple filters that determine whether to allow or block packets based on information in protocol headers.

static routes Manual entries in a router's routing table by an administrator.

stream cipher A type of encryption algorithm that encrypts one bit at a time.

strobe scan A type of port scan that probes ports commonly used by specific programs and attempts to use the programs. A common type of strobe scan looks for active Trojan programs that have already circumvented security measures and are running on the scanned system. If an attacker can find an installed Trojan, he can use it instead of having to install a new one.

stub network A network segment on a larger network's edge; a dead-end segment.

stub router A router that connects a stub network to the larger network.

subnet mask A value that defines which part of a computer's IP address is its network identifier and which part is the host identifier.

subpoena A legal document that requires a person to appear in court, provide testimony, or cooperate with law enforcement.

substitution box (S-box) function A cryptographic primitive that transforms a number of input bits into a number of output bits and produces a fixed or dynamic lookup table.

supernetting See *route summarization*.

survivability The ability to continue functioning in the presence of attacks or disasters.

Survivable Network Analysis (SNA) A security process that assumes a computer system will be attacked and follows a set of steps to build a system that can survive the attack.

symmetric algorithm A type of mathematical formula in which the key for encrypting cleartext is the same key for decrypting ciphertext.

syslog server A log management service.

target-to-console ratio The number of target computers on a network managed by a single command console.

Temporal Key Integrity Protocol (TKIP) An encryption method devised as a replacement for WEP in WPA. TKIP is based on the same RC4 mechanisms used in WEP, but it includes a method for generating new keys for each packet. It also incorporates automatic AP-client synchronization of unicast encryption keys, so keys do not have to be changed manually, as they do in WEP.

Terminal Access Controller Access Control System Plus (TACACS+) An authentication, authorization, and accounting service protocol.

Threat and Risk Assessment (TRA) An approach to risk analysis from the standpoint of threats and risks to an organization's assets and the consequences if those threats and risks occur.

threats Events and conditions that could potentially occur; their presence increases risk.

three-pronged firewall A firewall with separate interfaces connected to an untrusted network, a semitrusted network, and a trusted network.

ticket-granting server (TGS) The part of the KDC that creates and distributes session keys used by clients to access resources. See also *Key Distribution Center (KDC)*.

ticket-granting ticket (TGT) A digital token sent from the Authentication Server to the client. The client presents the TGT to the TGS to obtain a session key to access the resource. See also *Key Distribution Center (KDC)* and *ticket-granting server (TGS)*.

tiger teams Special teams assembled to test a network.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Tinkerbell program A program in which network connections are scanned and alerts are generated when logons are attempted from a suspicious IDPS or connection attempts are made from suspicious sites.

topology The way in which systems in a network are connected to one another.

Transmission Control Protocol/Internet Protocol (TCP/IP) A suite of protocols for transmitting information from point to point on a network.

Transport Layer Security (TLS) A protocol designed to provide additional security for Internet communication. TLS uses a hashed message authentication code (HMAC) to combine the hashing algorithm with a shared secret key. TLS splits input data in half, processes each half with a different hashing algorithm, and recombines them with an XOR function.

Triple DES (3DES) An enhanced variation of DES that uses three 64-bit keys to process data. See also *Data Encryption Standard (DES)*.

Trojan program A harmful computer program that appears to be something useful to deceive a user into installing it.

true negatives Legitimate communications that do not cause an IDPS to set off an alarm.

true positive A genuine attack detected successfully by an IDPS, in contrast to a true negative or false positive.

trust anchor The top-level digital certificate in a PKI chain.

tunnel The connection between two endpoints in a VPN.

tunnel method encryption A method of key exchange that encrypts both the header and data components of a packet and encapsulates it within a new packet that has a different header.

tunneling protocols Network protocols that encapsulate (wrap) one protocol or session inside another.

two-factor authentication Authentication that requires at least two forms of verification for a user to be granted access.

unicast A transmission in which one packet is sent from a server to each client that requests a file or application.

Unified Threat Management (UTM) A term for products that integrate a variety of security features into a single application, device, or product. These features include VPN and remote access services, firewalls, intrusion detection and prevention functions, and management consoles.

User Datagram Protocol (UDP) A core transport protocol of the TCP/IP suite. UDP is connectionless, meaning it does not ensure delivery or provide ordering, as TCP does. UDP is much faster, and is useful for transmissions that require speed over reliability. UDP relies on upper-level protocols for error-checking and sequencing services.

vanilla scan A type of port scan in which all ports from 0 to 65,535 are probed, one after another.

variable length subnet masking (VLSM) A means of allocating IP addressing according to the network's needs that involves applying masks of varying sizes to the same network. This method creates subnets within subnets and multiple divisions of an IP network.

virtual private network (VPN) A cost-effective way for networks to create a secure private connection using public lines (usually the Internet). VPN endpoints establish connections (tunnels) to transmit and receive data, and then tear down the connections when they are no longer needed. Combinations of encryption, authentication, and encapsulation help ensure the confidentiality, privacy, and integrity of information.

virtual terminal (VTY) Virtual session access points for simultaneous access to a Cisco device.

virus Computer code that copies itself from one place to another surreptitiously and performs actions that range from benign to harmful. Viruses require some user action, such as clicking an executable attachment or viewing an infected Web page, to enable them to launch.

virus scanning The process of examining files or messages for filenames, patterns, extensions, and other indications that a virus or other malware is present.

voltage standing wave ratio (VSWR) A type of loss caused by differences in the equipment itself. For example, if one part of the equipment has a different impedance than others, the signal can be reflected back within the device, and the reflected power can cause the device to burn out.

VPN appliance A hardware device designed to terminate VPNs and join networks.

VPN client A router or an OS that initiates a connection to a VPN server.

VPN domain A group of one or more computers that the VPN hardware and software handle as a single entity. This group uses the VPN to communicate with another domain.

VPN protocols Standardized communication settings that software and hardware use to encrypt data sent through a VPN.

VPN quarantine A method to address the problem of remote clients not meeting an organization's security standards. Quarantine places remote clients in a secured area while they are checked to ensure that software updates and current patches have been applied, antivirus software has been installed and updated, and other policies are in compliance.

VPN server A computer configured to accept VPN connections from clients.

vulnerabilities Situations or conditions that increase threats, which in turn increase risk.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

wardriving Driving around with a laptop and WNIC in RF monitor mode to detect unsecured wireless signals.

watt A measure of power or the rate at which work is done.

wave A means of transporting energy from one place to another without physical movement of material.

wave propagation The pattern of wave dispersal as waves travel from sending to receiving antennas.

wavelength The distance between two successive amplitude peaks.

Wi-Fi Protected Access (WPA) A subset of the 802.11i standard that addresses encryption and authentication; it uses IEEE 802.1x or preshared keys for authentication and uses TKIP for encryption. See also *IEEE 802.1x* and *Temporal Key Integrity Protocol (TKIP)*.

Wi-Fi Protected Access version 2 (WPA2) The encryption and authentication architecture based on the final ratified IEEE 802.11i standard. WPA2 uses preshared keys or IEEE 802.1x for authentication and uses AES for encryption. See also *Advanced Encryption Standard (AES)* and *IEEE 802.1x*.

Wi-Fi Protected Setup (WPS) A protocol from the Wi-Fi Alliance designed to automate key distribution in small office and home networks.

wireless networking Any exchange of data between computers and other devices that uses RF transmission methods instead of traditional wired cables.

worm Computer files that copy themselves repeatedly and consume disk space or other resources. Worms do not require user intervention to be launched; they are self-propagating.

worst-case scenarios Descriptions of the worst consequences that can befall an organization if a threat occurs.

X.509 An International Telecommunication Union standard for PKI that specifies standard formats for public key certificates, a strict hierarchical system for CAs issuing certificates, and standards for certificate revocation lists. X.509 certificates use RSA for key generation and encryption, and MD5 hashes to verify the certificate's integrity.

zone transfer The communication of a zone file from the primary DNS server to secondary DNS servers for updating.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Index

3DES encryption, 174
802.11 standards, 214–215, 236–237

A

AAA logging, 133, 145
AAA security, 136–137
abnormal network traffic signals, 95–96
absorption, RF signals and, 198, 217
acceptable use policy, 501–502, 512
access control authentication, 457
access control entry (ACE), 136, 145
access control list (ACLs), 14, 459
  create and use, 128–129
  creating, 156
  defined, 145
  extended ACLs, 131
  named ACLs, 132–133
  standard, 130
access point, 197
  defined, 217
  researching, 225–226
  wireless, 206–208
access point mode, remote wireless bridge, 210–211
accountability, 289, 290
account management, IDPS, 535
acknowledgment (ACK) frame, 255
  control frame, types, 230
acknowledgment flag (ACK), 89
acknowledgment number, TCP elements in packets, 88
ACK scans, 100
ACL. See access control list
acoustic cryptanalysis, side channel, 180, 401
active attacks, 233, 255
  common, 233–234 (table)
active defense-in-depth, 536–537, 541
Active Directory, 457, 458
active scanning, 232–233, 255
Active Server Pages (ASP), 446
ActiveX control attacks, 452, 461
adaptation and evolution, Survivable Network Analysis (SNA), method, 484
address
  assigning IP to interfaces, 151–152
  falsified IP, 101
  MAC, filtering, spoofing, 240
  resolution protocol processes, 121
"Address Allocation for Private Internets," 101
addressing
  classes, 38
  IPv6, 62
  private IP ranges, 38–39
  TC/IP and, 37–38
  unicast, multicast, anycast, IPv6 and, 62–63
Address Resolution Protocol (ARP), 55, 145
Adeyemi, Adeniyi, 4
ad-hoc mode, 236
Advanced Encryption Standard (AES), 164, 174, 184, 246, 255, 400–402
agents, knowledge-based IDPS and, 270–271, 290
Aircrack-ng, 233, 250
AirSnort, 250
alarm, alert message, IDPS response, 287
alert messages, IDPS send, 287
alert settings, IDPS and, 272
Altamirano, Luis Robert, 4
always-on connectivity, 9
amplitude, 217
amplitude key shifting (ASK), narrow band transmission, 202, 217
amplitude modulation, 201, 217
analog modulation, 201
analog RF signal, 200, 217
analyzing packet attacks, 85–88
anomaly-based IDPS, 268
anomaly detection systems, 267–269, 290
Anonymous, online activists, 3
Ansir Khan, 4
antennas, 208–210, 217
antispoofing logging, 135–136, 145
antivirus protection, layered defense, 13
anycast address, 461
  IPv6 and, 63, 67
  root servers and, 440
Apache Web Server
  configuring security settings in, 458–459
  installing, 378–379
AP logging functions, 250–251
AP placement, planning, 242–243
application vulnerability, 8 (table)
ARP poisoning, 248
ARP tables, 121, 145
assessments, risk and security, 249
asset(s)
  defined, 479, 512
  electronic, 491–492
  network, 490
  types of, 480–481
  valuation, 486
asset definition, 485
asset identification, conducting, 520–521
association
  defined, 255–256
  with wireless network, secure WLAN and, 235–236
association flood (DoS attack), 233
association request, management frame type, 230
asymmetric algorithms, 163, 184
asymmetric cryptography system, 170–172
Asynchronous Transfer Mode (ATM), 137
atomic attack, 85
attacks
  atomic attack, 85
  database of signatures or behaviors, 273–274
  denial of service attacks (DoS), 233, 441–443
  discover Web-based, 474
  distributed denial of service attack (DDos), 3, 7 (table), 21
  presentation on wireless methods of, 261
  scanning, wireless networking and, 232–234
  side channel, 401
  single-packet attacks, 107
  single-session attacks, 270
  SQL injection attacks, 445, 460.

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

attacks, against Web users
ActiveX control, 452
file attachment, 451–452
phishing, 449–451
attack techniques, Web site, 443–453. *See also* Web site attack techniques
attenuation. *See* loss
auditing, 17, 20, 512
compliance, with VPN, 419–320
security. *See* security auditing
tools, 250
wireless, 249
authentication
802.1x, 244
authorization, and accounting (AAA) logging, 133, 145
by assertion, 405
basic layered defense and, 13, 20
on Cisco router. *See* Cisco authentication/authorization
hardening Windows Web servers, 456–457
header (AH), 177–178, 396–397, 422
IPv6 core protocol 57–58
IPsec 396–397
Kerberos, 405–406
logging VPN activity, 419
management frame type, 230
security, layered defense, 12–13, 20
virtual private network and, 404–406
wireless, 236–239
authorization, Cisco router. *See* Cisco authentication/authorization
Automatic Private IP addressing (APIPA), 62
auxiliary (AUX) port, 122, 145
AUX passwords, 138
availability, 11, 20

B

backbone
Internet communication, routers and, 439
Internet service providers, 438
network of NSP, 439-440 (figure)
back doors, 106
random scans of, signature analysis and, 97
threats to network security, 5, 20
Back Orifice 2000, Trojan, 98
backups
IDPS security device, 534–535
securing information, 492
bandwidth
managing, 538–539
throttling, 278
Bank of New York Mellon, 4
banners
defined, 145
router log on, 139–140
Barracuda NG Firewall, 311
basic authentication, layered defense, 13, 20
Basic Service Area (BSA), 212, 217
Basic Service Set (BSS), WLAN configurations, 212, 217
bastion host
backups and auditing, 364
choosing, 358–359
defined, 373
general requirements to create, 359
hardening, 360–361
honeypots, using, 362–363
memory and processor speed, 360
network location, 360
operating system, choose, 359–360
select host machine, 359
services, selecting, 361–362
user accounts, using, 363
beacon, management frame type, 230
Beck-Tews exploitation technique, 248
behaviors, database of, 273–274
binary numbers, 37
biometrics, 13, 21
bit shuffling permutation functions, 160
black box, server attacks and, 443
blacklist, IDPS and, 272
block cipher, 184
encryption algorithms, 163
integral cryptanalysis, 183
Rijndael, 164
symmetric, 160
Twofish, 164
XSL attacks, 181
block Web sites, 383–384
Blowfish block cipher, 163–164
bookmarks, security-related
assembling for RSS feeds, 546
assembling for Web pages, 544–546
botnets, 461
Internet host security and, 442–443
threats to network security, 5, 21
bounds checking, 444
BOYD (bring your own device), 504
branch offices
multiple firewalls, 353
protecting with multiple firewalls, 350–351
broadcast, 67
broadcasting, TCP/IP and, 43
brute force password guessing, 8, 234
BSSID
data frame fields, 232
management frame field, 229
buffered logging, 145
buffer overflow attacks, 443–445, 460
buffOne variable, 444–445
BugTraq, 540
burglary, response to, 509
business needs, VPNs and, 389–391

C

CA (Computer Associates) Internet Security Suite, 309
cache poisoning, DNS, 442
Camellia block algorithm, 181
Candidate vulnerability status, 83
Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA), 234
carrier signal, 200, 217
cell phone security, 252–253. *See also* mobile device security
Center for Internet Security (CIS), 458–459, 550 (table), 539
Central Intelligence Agency, 3
centralized data collection, 528–529, 541
CERT Coordination Center, 143–144, 539, 550 (table)
certificate, asymmetric cryptography system, 171
certificate revocation lists (CRLs), asymmetric cryptography system, 171

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

certification authorities (CAs), asymmetric cryptography system, 171, 459
certifications, 540
Certified Wireless Network Professional organization, 551
challenge-based authentication, 456
challenge/response authentication, layered defense, 13, 21
change management, security event management, 531–532
Check Point Firewall, 14
Check Point NGX, 308, 310
checksum, 84, 106
  ICMP header element, 86
  TCP elements in packets, 88
chipping code, 205, 217
chosen ciphertext attack, 180–181
chosen plaintext attacks, 180–181
cipher command, using, 191–192
ciphertext, 158, 184
Cisco ASA series, firewall, 310
  configure, 367–372
Cisco authentication/authorization, 136–137
Cisco Internetwork Operating System (IOS), 121–122, 145
Cisco PIX line, 308
Cisco router logging, 133
Cisco's Layer 2 Forwarding (L2F), 393
C language, 444
Classless Interdomain Routing (CIDR), 12, 42–43, 67, 285
cleanup rule, 320, 328, 330
clear to send (CTS) frame, 256
  control frame, types, 230
client security, ensuring, 416–417
client-to-site VPN, 388, 422
co-channel interference, 197, 212, 217
code
  buffer attacks and, 443–445
  ICMP header element, 86
  malicious, threats to network security, 5
code analysis, IDPS response, 287
command console, IDPS interface, 273, 290
commercial proxy servers, 358
Common Gateway Interface (CGI), 457
  attacks, 104, 106
Common Vulnerabilities and Exposures (CVE)
  detecting attacks with, 81 (figure)
  how it works, 80–82
  scanning descriptions, 82–83
  table, 549
Comodo Internet Security, 8
CompTIA, 540, 551
computer emergency response team (CERT), 507. *See also* CERT
Computer Security Resource Center (CSRC), 550 (table)
Computing Technology Industry Association, 551
Conficker virus, 5
confidentiality, 11, 21
configuration
  IPv6, 63
  management, router security and, 144
connectivity, always-on, 9–10
consequences, risk assessment and
  describing, 486 (figure)
  risk analysis and, 482–483
console (CON) port, 122, 145
console passwords, 138
control frame
  defined, 230, 256
  fields in, 231
  types, 230–231
control Internet services, 325–329
controlling dynamic content, 475
control objects, ActiveX attacks and, 452
convergence, 124, 145
corporate phishing e-mail, 450 (figure)
cost-benefit analysis, 483, 512
cost estimates, Project Risk Analysis, 487, 488 (figures)
coWPatty, 233
crackers, 3, 21
crime rates, as threat, 481
criminal computer offenses, 496–497
cryptanalysis
  chosen ciphertext, chosen plaintext attacks, 180–181
  defined, 184
  differential, 182–183
  integral cryptanalysis, 182
  passive attacks, 180
  random number generator, 181
  related key attacks, 181–182
  side channel attacks, 179–180
  XSL attacks, 181
cryptographically secure pseudorandom number generator (CSPRNG), 160
cryptographic primitives, 158–159, 184
  substitution box functions, 160
cryptographic protocols
  about, 158
  Blowfish, 163–164
  defined, 184
  digital signatures, 168–169
  encryption algorithms, 162–163
  exclusive OR function, 159–160
  Feistel network, 160
  hashing algorithms, 165
  hashing functions, 162
  key management, 169
  key size in encryption algorithms, 163
  Message Authentication Code, 167–168
  Message Digest 5, 165–166
  permutation functions, bit shuffling, 160
  primitives, 158–159
  private key exchange, 169–170
  public key cryptography standards, 172
  public key exchange, 170–171
  Rijndael, 164
  Rivest Cipher, 164
  Secure Hash Algorithm, 166
  Twofish, 164
  X.509, 172–173
cryptography, 184
cryptography standards
  about, 173
  advanced encryption standard, 174
  authentication standards, 177–178
  data encryption standards, 173
  encapsulating payload, 178–179
  Internet Protocol Security (IPsec), 175–176
  Internet and Web standards, 174
  IPsec, about, 175–176
  IPsec, components, 176–177
  Secure Shell, 174–175

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Secure Sockets Layer, 175
Transport Layer Security (TLS), 175
triple DES, 173–174
CSMA/Collision Detection (CD), 234
CyberEYE program, 84
CyberSecurity Watch, 3
Cyclic Redundancy Check (CRC), 43, 162

D

Daemen, Joan, 164
data, ICMP header element, 87
data assets, 480
database, of attack signatures/behaviors, 273–274
database attacks, SQL injection and, 445
database-generated error message, 447 (figure)
data collection
centralized, 528–529
distributed, 529–530
data confidentiality, 457
data encapsulation, 492
data encryption, 491. *See also* encryption
Data Encryption Standard (DES), 173, 184, 399
data frames
defined, 231, 256
fields in, 232
datagram, 67
data transfers, encrypting, 491
dBm (decibel-milliwatt), 199, 217
deauthentication, management frame type, 230
decibels, 199, 217
default key, 239, 256
default routes, 122, 145
default settings, changing hardware/software, 243–244
default values, wireless devices, 260
defense, impact of, 19
defense, layered
about, 11–12
authentication, password security, 12–13
physical security, 12
"Defense-in-depth: Security for Network-Centric Warfare" (Galik), 537–538
defense-in-depth (DiD), 12, 21, 541
active, 536–537
improve, 536
security layers, adding, 537–538
defense zones, 537–538
degaussing, 539, 541
demilitarized zone (DMZ), network security, 15, 21, 316
denial of service attacks (DoS), 233, 441–443 *See also* attacks
denial of service traffic signal, 96
Department of Homeland Security, 82
deprecated vulnerability status, 83
desktop computer WNICs, 207
destination address, 32-bit IP address, 414
destination and source address
data frame fields, 232
management frame field, 229
destination IP address
IP headers and, 45
IPv6 core protocol, 57
destination options header, IPv6 core protocol, 57
destination port, TCP elements in packets, 88
detection, IDPS servers, 272–273
detection methodologies
anomaly and signature detection systems, 267–269
stateful protocol analysis, 269–270
detection methods, combining IDPS, 283
detection-related code, 272
dictionary, attack, 234
differential cryptanalysis, 182–183
Differentiated Services Code Point (DSCP), IP header structure, 45
Diffie-Hellman encryption, 395
diffraction, 217
diffused IR transmission, 195, 217
digest authentication, 459
digital certificate authentication, 456
digital modulation, 201–203
digital RF signals, 200, 217
digital signatures, 168–169, 184
directed broadcasts, 43
directed IR transmission, 195, 217
direct sequence spread spectrum (DSSS), 203–204
disassociation, management frame type, 230
discrete logarithms, 160
discretionary access control (DAC), 19, 21
distance-vector routing protocols, 124, 146
distributed data collection, 529–530, 541
distributed denial of service attack (DDos), 3, 7 (table), 21. *See also* attacks
DMZ
protecting with multiple firewalls, 350
firewall configurations, multiple, 348–349
DNS. *See* Domain Name System (DNS)
DNSSEC
attack protection, 455–456
using, 192
DNS Security Extensions (DNSSEC), 455–456, 461
domain controller, security and, 458
domain group, creating for VPN remote access, 426–427
domain information, finding, 464–465
Domain Name System (DNS), 67, 439–440, 460, 461
cache poisoning, 442
hierarchy, 441 (figure)
information leakage, 442
Internet Protocol Version 4 (IPv4) and, 53-54
resolution, packet filtering and, 325–326
servers, hardening, 453–455
spoofing, 442 *See also* DNS
DoS attacks, 233, 441, 442–443
download/install router simulation program, 149–150
drop, IDPS response, 287
DShield, 549 (table)
dual-homed hosts, 344–346, 353, 373
duration
control frame fields, 231
data frame fields, 232
management frame field, 229

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

dynamic application layer protocol analysis, stateful protocol analysis, 269
Dynamic Host Configuration Protocol (DHCP), 37, 55
dynamic routes, 124, 146
Dynamic WEP, 239

E

EAP-Transport Layer Security (EAP-TLS), 245
EAP-Tunneled TLS (EAP-TTLS), 245
economic impact, analyzing for risk analysis, 487–489
EIRP, 208
electromagnetic radiation, 194–195, 217
electronic assets, 491, 512
e-mail
configuration, packet filtering, 326–327
corporate phishing, 450 (figure)
file attachment attacks, 451–452
gateways, 18
phishing attacks and, 449–451
security policy and, 504
threats to network security, 8
employees
disgruntled, threats to network security, 3–4
security breaches and, 11
EM spectrum, 194, 218
Enable passwords, 137
enable secret passwords, 137–138
encapsulation, VPNs and, 393–394, 422
encapsulating security payload, 178–179, 397–398, 422
Encapsulating Security Payload (ESP) header, IPv6 core protocol, 58
encryption
authentication and, strong, 244
configuring in Word 2010, 192
Diffie-Hellman, 395
nonrepudiation and, 11
password, 137, 138–139
risk management and, 491
schemes, VPNs and, 399–404
standard for 802.11i, 246
using TrueCrypt, 189–190
VPNs and, 398–399
encryption algorithms, 162–163, 184
end users, define roles of, 505
Enhanced Interior Gateway Routing Protocol, 125
enterprise firewall, 310
Entry vulnerability status, 83
Equivalent Isotropically Radiated Power (EIRP), 200, 218
error messages, 444 (figure)
database-generated, 447 (figure)
Internet Protocol Version 6 (IPv6), 58–59
eSafe Content Security, 308
escalated, 290
escalation procedures, 288–289, 508, 512
mapping, 509
ethical hackers, 4, 249
event horizon, 269
event monitoring, 526–528, 541
exclusive OR function, 159–160, 184
Explicit Congestion Notification (ECN), IP header structure, 45, 47
exploitation of rogue devices, wardriving and, 234–235
exploits, application vulnerability, 8 (table)
exposure, 481, 512
extended IP ACLs, 131–132
Extended Protection, 457
Extended Service Set (ESS), 212, 213, 218
Extensible Authentication Protocol (EAP), 244, 256
extranet
defined, 502, 513
third party access, security policy and, 502

F

facility
age of, as threat, 481
security, 506
fading, 218
failover firewall, 350, 373
FakeAP, 250
false negatives, 267–269, 288, 291
false positives, 267–269, 288, 291
falsified IP address, 101
falsified port number, 102
Family Educational Rights and Privacy Act (FERPA), 11, 253
Fannie Mae, 3–4
fast basic service set transition (FT), 218
fault tolerance, Survivable Network Analysis (SNA), method, 484, 513
Federal Communication Commission (FCC), 215, 218, 222–223
Federal Information Processing Standard (FIPS), 173
Feistel network, 160, 173, 184
Fibre Distributed Data Interface (FDDI), 137
file attachment attacks, 451–452
file system monitoring, IDPS response, 287
File Transfer Protocol (FTP), 37, 438
analyzing signatures, 91–93
filtering
IP addresses, 320
MAC address spoofing and, 240
by TCP/IP and UDP ports, 321–323
by status, 323–325
filtering content, proxy servers and, 358
filter rules, developing IDPS, 284–285
fingerprints, 504
finished flag (FIN), 90
FIN packets, 100–101
FIN scans, 100
firewall, using VPNs with, 411–414
firewall appliances, 308
firewall configurations
about, 344
branch offices, protecting with multiple, 350–351
choose a, 353
DMZ firewall configurations, multiple, 348–349, 353
dual-homed hosts, 344–346, 353
example, 367–372
multiple, 350
protecting DMZ with multiple firewalls, 350
reverse firewalls, 351–352, 353
screened hosts, 346–347, 353
screened subnet DMZs, 347–348, 353
screening routers, 344, 353
firewall policy, 318–319, 330
firewall programs
advantages/disadvantages, 311 (table)

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

antivirus protection and, 13
application traffic policy, 318–319
commercial software, 309–310
defined, 306–308
evading, IDPS and, 118
filtering
by TCP/IP and UDP ports, 321–323
by service, 323–325
IP addresses, 320
free, 309
hands-on project, 25–27
hardware, 310–311
IPsec, enforcing policies, 334–338
network security, 14
packet filtering, position and, 314–317
permissive *vs.* restrictive policies, 14–15
proxy servers that include, 358
rule base
designing, 339–340
improving, 340–341
simplicity, 319–320
security policy rule set, 318
software, recommending, 341
software-based, 309–310
stateful packet filtering, 313–314
stateless packet filtering, 312–213
Telnet rule, 338–339
threats to network security, 8
Windows firewall, advanced settings, 332–334
firewall selection parameter, 384
Firewalls Mailing List, 540
fixed-mobile access, 226
flags
identify, 89–90
illegal TCP, 102–103
IP headers and, 45
TCP control, 323–324
TCP elements in packets, 88
TCP headers and, 47–48
flash drives, encrypting, 491
Flexible Authentication via Secure Tunneling (FAST), 246
flooded broadcasts, 43
flow label, IPv6 core protocol, 56
footer, 43, 67
forged deauthorization, 233
forged disassociation (DoS attack), 233
Fortinet FortiGate series, firewall, 311
Forum of Incident Response and Security Teams (FIRST), 550 (table)
Fourth Amendment, U.S. Constitution, 496–497, 513
fragmentation, 67
Fragmentation abuses, suspicious events, 102–104
fragment header, IPv6 core protocol, 57
Fragment Offset, IP header and, 45
frame body
data frame fields, 232
management frame field, 229
frame check sequence
control frame fields, 231
data frame fields, 232
management frame field, 229
frame control
control frame fields, 231
data frame fields, 232
management frame field, 229
FreeBSD Project, 360
free space path loss, 209
freeware proxy servers, 357–358
frequency, 194–195, 218
frequency division multiplexing (FDM), 203, 218
frequency hopping spread spectrum (FHSS), 204, 218
frequency modulation (FM), 201, 218
frequency shift keying (FSK), 202, 218
Fresnel zone, 209, 218
FTP server, installing, 113–115
FTP traffic, capture with TShark, 116
FTP transactions, packet filtering, 327
fully qualified domain names (FQDNs), 54, 67, 325–326
functional managers, define roles of, 505

G

gain, 197–198, 218
Galik, Dan (Captain USN), 537–538
Gateway 2010, installing, 375–378
gateways, wireless, 211–212
gateway-to-gateway VPN, 388
Generic Routing Encapsulation (GRE), 393–394, 422
Geotrust, 459
Global Information Assurance Certification (GIAC), 551
Global Positioning System (GPS), 253
global unicast address, IPv6 and, 62
Gnu Datagram Transport Layer Security (TLS), 83
GoCertify, 540
government operations, threats to network security, 4
Gramm-Leach-Bliley Act, 11
Guide to Intrusion Detection and Prevention Systems (IDPS), 270

H

hackers, threats to network security, 2–3, 21
hacktivists, 3, 21
Halvorsen, Michael, 248
handheld devices, researching security for, 260–261
handshake mechanism, 529
hardening
DNS servers, 453–455
router, 143–144
Windows OS, 457
hardening Windows Web servers
access control, 457
authentication, 456–457
controlling dynamic content, 457
data confidentiality, 457
security considerations, 457–458
shared configuration, 457
hardware
change/add, 535
firewalls, 310–311
inventory, conduct, risk management and, 490
securing, risk management and, 489–490
VPS, 388–389
hardware/software
changing default settings, 243–244
inventory, collecting, 518–520

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

hashed message authentication code (HMAC), 175
hashing algorithms, 165, 184
hashing functions, 162, 184
hashing message authentication (HMAC), 403
header(s)
bad information, 83–84
defined, 67
ICMP, elements, packet analysis, 86–87
IPv6 core protocol, 56–57
header checksum, IP header and, 45
header length, TCP elements in packets, 88
Health Insurance Portability and Accountability Act (HIPAA), 11, 253
highly directional antenna, 210
honey pots, using, 362–363
Hop-by-Hop Options header, 57
hop limit, IPv6 core protocol, 57
hopping code, 205, 218
host-based configurations, 270
host-based IDPS(HIDPS), 291
about, 279
comparing with NIDPS, 282
configuring, 279–281
host, choosing, 281–282
host identifier, 37, 67
hybrid configuration, 408, 422
hybrid IDPS, 291
about, 282
advantages and disadvantages of, 283
detection methods, combining IDPS, 283
sensor locations, combining IDPS, 283
hybrid topography, VPNs and, 408–409
HyperTerminal, 138
Hypertext Transport Protocol (HTTP), 8, 37, 438

I

IANA, protocol number assignment, 102
IBM, Data Encryption Standard (DES), 173
IBM Internet Security Systems, 310
ICMP
destination unreachable message, 49
ID numbers, 87, 106
message types, packet filtering, 327–329
packets, use TShark to capture, 111–113
v4 redirect messages, 55
v4 router discovery, 55
identification, IP headers and, 45
identifier (BE) ICMP header element, 87
identifier (LE), ICMP header element, 87
IDPS. *See* intrusion detection and prevention system
IDPS management server, 271–272, 291, 541
backups, 534–535
change/add software, 535
hardware, change/add, 535
maintain system, 534
manage accounts, 535
rule management, 535
user management, 535
IDPS signatures, evaluating, 530–531
IEEE 802.11, 256
media access control, frames, 228–332
wireless networking standard, 213–214, 218
authentication, 256
IEEE Computer Society, 550 (table)
IEEE 802.1x standard, 244, 256
independent auditing, 534, 541
Independent Basic Service Set (IBSS), 213, 218
Industrial, Scientific, and Medical (ISM), 198
industry, security threat and, 481
information leakage, DNS, 442
informational messages, Internet Protocol Version 6 (IPv6), 58–59
informational traffic signature, 96
infrared (IR) radiation, 218
infrared transmissions, 195, 219
initialization vector (IV), 238, 256
inline firewalling, 278
inline only sensor, 278
inline sensor, 276, 278, 291
inSSIDer, installing, 223–225
Institute of Electrical and Electronics Engineers (IEEE), 212
integer factorization, 160
integral cryptanalysis, 182
integrity, 11, 21
intentional radiator, 208, 219
interface modules, 137
interference radio frequencies, 197
Intermediate System-to-Intermediate System, 125
internal network, from remote VPN client, accessing, 429–431
International Information Systems Security Certification Consortium (ISC2), 540, 551
International Telecommunication Union standard, for PKI, 172–173
Internet, 46
secure use of, 504
Web standards, cryptography and, 174
Internet, structure of
DNS security, 442
Domain Name System (DNS), 439–440
Internet service providers, 439
network access points, 439
routing security, 442
tier system, 439
understanding, 438
weak points in, 440–441
Web site attack techniques, 443–453. *See also* Web site attack techniques
Internet Assigned Numbers Authority (IANA), 415, 550 (table)
Internet communication backbone, routers and, 439
Internet Control Message Protocol (ICMP), 46–47, 67
message abuse, 7 (table)
Internet Control Message Protocol v6, firewall programs and, 324
Internet Engineering Task Force (IETF), 172, 184, 394–395, 422, 550 (table)
Internet Explorer
Java applet attacks and, 453
security settings, 465–468
Internet Header Length (IHL), 56
Internet host security, 442–443

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Internet Information Services (IIS), 456
configure security for, 470–473
installing, 468–469
Internet Key Exchange (IKE), 176, 395, 422
Internet Message Access Protocol (IMAP), 37
Internet Protocol Security (IPsec), 175–176, 184, 394–395, 422. *See also* IPsec
Internet Protocol Version 4 (IPv4)
defined, 67
domain name system and, 53–54
flags in header, 324
fragmentation analysis, 10–104
ICMP messages, 46–47
IP datagrams, 43
IP header structure, 44–46
IPv6 *vs.*, 77
messages, packet filtering, 327–329
packet fragmentation, 49–51
TCP headers and, 47–48
TCP life cycle, TCP three-way handshake, 51–53
UDP headers and, 48–49
use Wireshark to capture pings in, 71–73
Internet Protocol Version 4 (IPv4), 37
Internet Protocol Version 6 (IPv6)
about, 54–55
addressing, 62
configuration, 63
control message protocol for, 58–59
core protocols, 55–58
defined, 67
exploring Windows implementation with Netsh, 154–156
firewalls and, 324
fragmentation analysis, 104
ipconfig, 63–64
IPv4 *vs.*, 77
messages, packet filtering and, 327–329
multicast listener discovery, 59–60
Netsh and, 65
netstat, 64–65
routing, 127–128
tunneling, 156
unicast, multicast, anycast addressing, 62–63
utilities, 63
Wireshark, examine ping packets with, 76–77
Internet security, threats to network security, 8
Internet Security Association Key Management Protocol (ISAKMP), 176, 395, 422
Internet Server Application Programming Interface (ISAPI), 457
Internet service provider (ISP), 439
Internet Storm Center, 549 (table)
Internet use policy, 504, 513
interval, suspicious attacks and, 100–101
intrusion, 291
intrusion detection and prevention system (IDPS), 15–16, 21
administrator assess damage, 287–288
anomaly in signature detection systems, 267–269
best practices, securing components, 283–284
components of, 270–274
defined, 291
escalation procedures, 288–289
evaluating IDPS products, 289–290
filter rules, 284–285
gathering data, 286
goals of, 266–267
host-based IDPS, 279–282. *See also* host-based IDPS
hybrid IDPS, 282–283
inline *vs.* passive, 303
installing IDPS database, 285–286
logging and reviewing events, 289
network-based IDPS (NIDPS), 275–279. *See also* network-based IDPS (NIDPS)
response, 287
sending alert messages, 287
intrusion prevention, 16
inverse masks, 130, 146
IOS command-line interface, working at, 150–151
IP
address, assign to interfaces, 151–152
address classes, 38
falsified address, 101
private address ranges, 38–39
ipconfig, IPv6 and, 63–64
IP datagrams, 43
IPDS load balancer, 277
IP packet header fields, VPNs and, 414–415
IP packet reassembly, stateful protocol analysis, 269
IPsec, 55
about, 175–176, 394–395
components, 176–177, 395–398
driver, 176, 395–396, 422
filters, 416
policies, enforcing, 334–338
IPsecurity Policy Management, 176. *See also* IPsec
IP spoofing, 101, 106, 441–442

J

jamming (DoS attacks), 233
Java
ActiveX control attacks and, 452
applet attacks, 453, 461

K

Kaspersky Internet Security, 310
Katmar Software, Project Risk Analysis, 487–489
Kerberos, 136–137, 422, 457
Kerberos authentication, 405–406
key derivation (generation), PRNGs and, 161
Key Distribution Center (KDC), 405, 422
key management, 169, 185
keystream, generate, 164
Kismet, 250
knowledge-based IDPS, 268

L

land attack, falsified IP address and, 101
LAN environment, security policy and, 505–506
LAN routing protocols, 125
LANs and their components. *See* Wireless LANs, components
layers, defense. *See* defense, layered, 12

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Layer 2 Tunneling Protocol (L2TP), 393, 422
legacy authentication, 457
Lightweight EAP (LEAP), 245
line-of-sight (LOS), 209
link-local unicast address, IPv6 and, 62
link-state routing protocols, 124, 146
Linux operating system
  packet filtering and, 14, 357 (table)
  Squid for, 358
load balancing
  firewalls for, 351, 373
  software, 350
local administrators, define roles of, 505
log file
  analysis, 17
  configuring, 17–18
logging, 146
  antispoofing, 135–136
  AP functions, 250–251
  buffered, 135
  examining Cisco, 133
  levels, 133–134
  options, 134–135
  reviewing events and IDPS, 289
  router security and, 143–144
  VPN activity, 419
logical access controls, security policy and, 502–503
logic bomb, threats to network security, 3–4, 21
login pages, SQL injection attacks and, 445–448
login redirection-based authentication, 456–457
loss, 198, 219
L2TP Access Concentrator (LAC), 393
L2TP filters, VPNs and, 416
L2TP Network Server (LNS), 393

M

MAC
  address filtering, spoofing, 240
  address spoofing, 233 *See also* media access control
macro virus, threats to network security, 6, 21
MAC tag, 167–168
mailing lists, security and, 539–540
malformed packet signatures, 102=103
malware signature, 84
ManageEngine's Firewall Analyzer, 308
management frames, 229, 256
management server, IDPS, 271–272
management, security policies and, 497–498
mandatory access control (MAC), 19, 21
man-in-the-middle attack, 7 (table), 235, 236 (figure), 457
many-to-one NAT, 365–367, 373
Massachusetts Institute of Technology (MIT), 405
maximum transmission unit (MTU), 57, 106
McAfee Network User Behavior Analysis, 272–273
Media Access Control (MAC), 55
Melissa macro virus, 6
memory, managing, 538
mesh configuration, 406–408, 422
mesh topography, VPNs and, 406–407
Message Authentication Code (MAC), 167–168, 185
Message Digest 5, 137, 165–166, 185
message digests, asymmetric cryptography system, 171
message filtering, 491
Message Integrity Check (MIC), 167–168, 246–249, 256
metrics, 125, 146
Microsoft company value statement, 476
Microsoft Forefront Threat Management Gateway (TMG), 358, 375–378
  configuring, 379–381
  configuring, as VPN server, 427–429
  exploring, 431–435
  features, 435
  logging features of, 381–382
Microsoft Word, 2010
  configuring encryption in, 192
  using a symmetric algorithm in, 188–189
Millimeter Wave, 198
mission statement, 476
misuse-based detection, 268
Mitre, Web site, vulnerabilities list, 82
M86 MailMarshal Secure Email Gateway, 8–9
mobile computers, encryption of, 491
mobile device security
  about, 252–253
  approaches to, 253–254
  researching, 260–261
mobility, wireless device portability, 240–241
ModSecurity module, 459
modulation, 201, 219
Monte Carlo simulation, 488, 513
mounted access, 285
M86 Security Labs, 442
multicast
  defined, 67
  IPv6, 63
multicasting, TCP/IP and, 43
Multicast Listener Discovery, Internet Protocol Version 6 (IPv6), 59–60, 67
multipath, 219
multipath distortion, 209
multiple entry point configuration, 410–411, 423
multiple firewall configurations, 350, 353
multiple-input multiple output (MIMO), 214–215, 219
multiple-packet attacks, 85, 107
multiple-session attacks, 270
multivariate quadratic equations, 401
mW, power measurement, 199
MyDoom worm, 5

N

named ACLs, 132–133, 146
narrowband transmissions, 202, 219
National Institute of Standards and Technology (NIST), 82, 174, 270, 400, 505, 550 (table)
National Science Foundation, NSFNET, 438
National Security Agency (NSA), 12
NAT *See* network address translation
needs assessment, system design, 477
Neighbor Discovery protocol, 55, 67
Neighbor Discovery protocol, Internet Protocol Version 6 (IPv6), 59–61
Nemesis, packet injection tool, 104
NetBios, 322
Netcom, 438
netfilter, firewall, 309

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Netscape, Java applet attacks and, 453
Netscape Communications Corporation, 175, 402
Netsh
exploring windows IPv6 implementation with, 154–156
IPv6 and, 65
Netstat, IPv6 and, 64–65
Netstat utility, 527
Netstumbler, 250
network access points (NAPs), 439, 461
Network Address Translation (NAT), 38, 55, 67
defined, 364–365, 373
many-to-one, 365–367
one-to-one, 365
network assets, 490, 513
network auditing, log files and
configuring log files, 17–18
log file analysis, 17
network-based IDPS (NIDPS), 291
capabilities, 278
comparing with HIDPS, 282
defined, 275
management, 279
positioning on network, 275–277
Network File System (NFS), 323
network identifier, 37, 67
Network Information System (NIS), 105
networking, wireless. *See* Wireless communication
network interface cards (NIC), 182
Network Mapper (Nmap), 99–100
network needs, to strengthen performance
bandwidth, 538–539
memory, 538
storage, 539
network security, goals of
antivirus protection, 13
auditing and log files, 17–18
confidentiality, integrity, availability, 11
connectivity, secure, 9–10
demilitarized zone (DMZ), 15
firewalls, 14
intrusion detection and prevention (IDPS), 15–16
nonrepudiation, encryption and, 11
operating system security, 13
packet filtering, 13–14
permissive/restrictive policy, 14–15
privacy, ensuring, 11
remote access, secure, 10
routing, access control and, 18–19
virtual private network, 16–17
network security, threats to
always-on connectivity, 9
common attacks and defenses, 6–8 (table)
disgruntled employees, 3–4
e-mail and communications, 8
government operations, 4
hackers, 2–3
Internet security, 8
malicious code
scripting, 8–9
social engineering, 6
terrorists, 4
viruses, worms, Trojan programs, 5–6
worms, 5
network security administrator, 453
network security policy, 504–505, 513
network sensors, IDPS components and, 270–271
Network Service Providers (NSPs), 438
routers, Internet communication backbones and, 439
network tap, 277
network traffic analysis, IDPS response, 287
network traffic filtering, IDPS response, 287
Network World Security Research Center, 550 (table)
newsgroups, security trends, 539–540
New Technology LAN Manager (NTLM), 457
next header, IPv6 core protocol, 57
NICs, wireless, 206
Nigerian money scams, 449
NIST Computer Security Division, 550 (table)
NIST Special Publication 800–94, 270
Nmap scans, explore signatures of, 117–118
nonces, PRNGs and, 160
nonrepudiation, 21, 168–169
encryption and, 11
non-return-to-zero (NRZ), digital modulation, 201
nonroot mode, remote wireless bridge, 211
Norton Internet Security, 309–310
Norton Security Suite Firewall, 306–307
null packets, 103, 107
null scans, 100

O

Oakley, IPsec, 176
Oakley protocol, 395
old school hackers, threats to network security, 3
omnidirectional antenna, 210
one-to-many communication, 63
one-to-one communication, 63
one-to-one NAT, 373
one-to-one network address, NAT and, 365
one-way functions, PRNGs and, 160–161
Open Shortest Path First (OSPF), 41, 125
open system authentication, 236, 237 (figure), 256
Open System Interconnections (OSI) Model
routing and, 120–121
TC/IIP protocols and, 36–37
operating system
choose for Bastion host machine, 359–360
mobile devices, 253
packet filters, 14
security, 13
operational auditing, 533–534, 541
options, IP headers and, 45
orthogonal frequency division multiplexing (OFDM), 204, 219
outbound Web access, 326 (table)

P

packet attacks
multiple, 85
single, 85
packet filtering, 13–14, 21, 442
adjust for VPNs, 414–416
based on position, 314–317

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

e-mail configuration, 326–327
FTP transactions, 327
rules, 323–325
stateless, 312–213
Web service rules, 325
packet fragmentation, 49–51
packet header discrepancies, 101
packet injection, attacks, 104, 107
packet monkeys, threats to network security, 3, 21
packet reconstruction, 355–356
packets, 17
analyzing, 85–88
packet sniffer, 85, 107, 250
pairwise master key (PMK), 246–249, 256
pairwise transient key (PTK), 246, 256
passive attacks, 180, 232, 257
passive only sensor, 278
passive scanning, 232, 257
passive sensors, 277, 278, 291
password, 492
AUX, 138
Cisco routers, 137
enable, 137
enable secret, 137–138
encrypting, 138–139
loss, response to, 509
protection, security policy and, 502–503
security, layered defense, 12–13
VTY, 138
patching, routers, 144
path obfuscation, attacks, 104
PATH variable, setting to work with TShark, 110–111
payload data, suspicious, 84
payload length, IPv6 core protocol, 56–57
Payment Card Industry Data Security Standard (PCI DSS), 11
PDAs, mobile security, 252–253. *See also* mobile device security
peer-to-peer network, 236
penetrating testing, 249, 250, 257
determining legal requirements of, 33–34
performance monitoring, 478–479
perimeter network, 15
permissive policy, network security and, 14–15, 21
permutation functions, bit shuffling, 160, 185
personal commercial firewalls, 309–310
personnel assets, 481
pharming, phishing attacks, 451, 461
phase modulation (PM), 201, 219
phase shift keying (PSK), narrowband transmission, 202, 219
phishing attacks, 449–451, 461
PHS flag, 102
physical assets, 480
physical security, 506
ensuring, wireless networking and, 242
examining for wireless devices, 260
physical security, layered defense, 12, 22
ping packets
analyze packet capture, 77
Wireshark, examine IPv4 with, 73–75
Wireshark, examine IPv6 with, 76–77
ping sweeps, traffic signal analysis and, 96, 107
pinhole vulnerabilities, 443
plaintext, 158, 185
point of presence (POP), 439, 462
point-to-multipoint configuration, 210–211
Point-to-Point Protocol (PPP), 9
Point-to-Point Tunneling Protocol (PPTP), 392–393, 423
point-to-point wireless bridging, 210–211
polarization, 208, 219
polar non-return-to-zero (polar NRZ), digital modulation, 201
policies
defined, 476
instituting secure wireless, 241–242
polymorphic buffer overflow attacks, 104
porous borders, 18
port(s)
defined, 22
filtering by, 321–323
malicious scanning, 7 (table)
number, falsified, 102
threats to network security and, 5
scanning tool, install, 27–28
Port 8080, 356
portability, wireless device, 240–241
Port Address Translation, 365–367
Portmapper, 105, 323
port scanning, 107
target system, 29–30
traffic signal analysis and, 97
Post Office Protocol (POP), 37
power, industry-standard measurements of, 199
power monitoring attacks, side channel, 180, 401
Power over Ethernet (PoE), 208, 219
power measurement, mW, 199
power-save poll (PS-Poll) frame, 257
control frame, types, 231
power supply, as threat, 481
PPP tunneling protocols, 393
PPTP filters VPNs and, 415–416
prevention capability, of IDPS server, 272–273
primitives, cryptographic, 151–159
privacy
employee right to, 497
ensuring, 11
private IP address ranges, 38–39
private key exchange, 169–170, 185
privileged access policy, 499, 513
PRNGs. *See* pseudorandom number generators
probabilities, 481, 513
defined, 481
sample threat, 482 (figure)
of threats, 481–482
probe request, management frame type, 230
probe response, management frame type, 230
profiles, 267–269, 291
Project Risk Analysis, Katmar Software, 487–489
Project Tinkerbell, 533
property loss/theft, 509
Protected EAP (PEAP), 245
protocol, IP headers and, 45
protocol ID, 414
protocol number assignment, 102
protocol state tracking, stateful protocol analysis, 269
Proventia security products, 310
proxy servers, 315–317, 330, 316
advantages/disadvantages, 357 (table)
about, 353–354

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

choosing, 356–357
commercial, 358
defined, 373
filtering content, 358
firewall functions and, 358
freeware, 357–358
goals of, 354–355
how they work, 355–356
parameters, 384
pseudorandom number generators (PRNGs), 160–163, 185
PSI Net, 438
Public DNS servers, 455
public key cryptography standard (PKCS) scheme, 165, 172, 185
public key exchange, 170–171, 185
Public-Key Infrastructure (PKI) system, 406, 456
public telecommunications infrastructure, 16–17
Push flag (PS), 89
PuTTY, 138

Q

Quality of Service (QoS) channels, 248
quarantine, 13
VPN, using, 418–419
Quote of the Day, 323

R

Rabin function, 160–161
radiation monitoring
side channel, 180
side channel attack, 401
radio frequency (RF)
FCC and, 215
measuring signals, 199–200
monitor mode, 232, 257
signal behavior, 197–198
signaling (RF), 200–201
transmissions, 196
waves, 194, 219
radio waves, 208
RADIUS, 136–137
random number generator attacks, 181
rank resources, risk management and, 490
reassociation request, management frame type, 230
reassociation response, management frame type, 230
receiver address, control frame fields, 231
recognition, Survivable Network Analysis (SNA), method, 484
reconnaissance traffic signature, 96
recovery, Survivable Network Analysis (SNA), method, 484
Red Hat Linux, 360
redundancy, securing information, 492
reflection, RF signals, 198, 219
refraction, RF frequency, 198, 219
registration authorities (RAs), asymmetric cryptography system, 171
related key attack, 181–182
remote access
secure, 10
secure shell, 140–141
security policy and, 503–504
Trojans (RATs), 84, 107
and wireless connection policy, 513
Remote Authentication Dial-In User Service (RADIUS), 133, 136–137, 146, 236
Remote Procedure Calls (RPC), 323
attacks, 7 (table)
defined, 107
suspicious events, attacks, 105
Remote Windows Shutdown, Trojan, 98
removable media, encryption of, 491
repeater mode, remote wireless bridge, 211
replacement costs, calculating, 516–518
Request for Comments (RFC), 38–39
request to send (RTS) frame, 257
control frame types, 230
research, summarizing, 262–263
reset, IDPS response, 287
reset flag (RST), 89
residual risk, 483, 513
resistance, Survivable Network Analysis (SNA), method, 484
response in, ICMP header element, 87
response teams, security policy, incident handling procedures, 507–508
restrictive policy, network security and, 14–15, 22
retina scan, 504
return on investment (ROI), 19, 498
return-to-zero, digital modulation, 201
reverse firewalls, 351–353, 373
review, security policy, 510
RF line of sight, 209
RF power gains, 199
Rijmen, Vincent, 164
Rijndael (AES), 181
Rijndael algorithm, 174, 185
Rijndael symmetric encryption algorithm, 400
risk, 479, 513
risk analysis, 511, 513
about, 479–480
assets, 480–481
asset valuation, 486
calculating, 486
conducting, 521–524
conduct ongoing, 492–493
consequences, describing, 486 (figure)
consequences of threats, 482–483
defined, 479
economic impacts, analyzing, 487–489
example, 500 (figure)
life cycle, 480
probabilities, 481–482
process of, 486
security controls, 483
security policy
develop from, 495–496
using to conduct, 499
software, 487–489
Survivable Network Analysis (SNA), method, 484
team meetings, 486
Threat and Risk Assessment (TRA), 485–486
vulnerabilities, 482, 486 *See also* risk management
risk assessment, conducting, 521–524
risk management, 513
defined, 489
encryption, using, 491
hardware inventory, conduct, 490
rank resources for protection, 490
securing hardware, 489–490
securing information, 491–492
Rivest, Ronald, 165
Rivest Cipher family, 164
Rivest Cipher v.4 (RC4), 238
Rivest, Shamir, Adelman, 164–165

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

RJ-45 jack, 206–208
roaming, 213, 219
rogue devices, 257
  defined, 234
  wardriving, exploitation of, 234–235
role-based access control (RBAC), 19, 22
role-based authentication, 503–504, 513
rollover cable, 122, 146
root mode, remote wireless bridge, 211
root servers, 440
routed protocol, 146
Router Advertisement messages, (IPv6 and), 63
router passwords, Cisco routers, 137
routers
  accessing, 121–122
  Internet communication backbone, 439
  packet filters and, 14, 442
  screening, 344
router security
  access control lists, 128–129
  antispoofing, 135–136
  AUX passwords, 138
  banner, 139–140
  buffered logging, 135
  Cisco router logging, 133
  console passwords, 138
  enable passwords, 137
  enable secret, 137–138
  enabling SSH on, 141–143
  encrypting, 138–139
  hardening, 143–144
  logging levels, 133–134
  logging options, 134–135
  passwords, 137
  remote access with secure shell, 140–141
  standard access control lists, 130
  use and rules, 129–130
  verifying SSH, 143
  VTY password, 138
router simulation program, downloading and installing, 149–150
route summarization, 125, 146
routing, 146
routing header, IPv6 core protocol, 57
Routing Information Protocol Version 1 (RIPv1), 41, 124
routing metrics, 124
routing process
  about, 120–121
  accessing routers, 121–122
  address resolution protocol processes, 121
  choosing protocol, 125
  distance-vector routing protocols, 124
  dynamic, 124
  of IPv6, 127–128
  link-state routing protocol, 124
  route summarization, 125–127
  routing metrics, 124
  routing tables, 122
  static routes, 122–123
routing protocol, 439
  choosing, 125
  defined, 146
Routing and Remote Access Services (RRAS), 419
routing security, 442
routing tables, 122, 146, 439
RPCbind, 105
RPC
  dump, 105
  NFS sweep, 105
  set spoof, 105
RSA key encryption, 164–165
RSA labs, 172
rule base, firewall/packet filtering, 317–320, 325, 330
  designing, 339–340
  improving, 340–341
rule header, 284
rule management, IDPS, 535
rule option, 284

S

SafeNet, 308
salts, PRNGs and, 161
sanitized, protected computers, 445
SANS Computer Security Newsletters and Digests, 550 (table)
SANS Institute, 360, 539
SANS NewsBites, 540
Sarbanes-Oxley Act (SOX), 11, 253
S-box function, 160
scanner, 107
scanning, attacks, wireless networking and, 232–234
scan throttling, 101, 107
scattering, RF signals, 198, 219
scopes, 62–64, 67
screened host, 346–347, 353, 373
screened subnet DMZs, 347–348, 353
screening routers, 344, 353
scripting, threats to network security, 8–9
script kiddie, threats to network security, 3, 22, 445
scripts, 22
search warrant, 497, 513
Secure Hash Algorithm (SHA), 166, 185
Secure Shell (SSH), 146
  cryptography standards, 174–175
  enable on router, 141–143
  remote access, 140–141
  verify, 143
Secure Sockets Layer (SSL), 164, 402, 423, 459
  cryptography standards, 175
secure WLAN implementation
  association with wireless network, 235–236
  default keys, 239
  MAC address filtering, spoofing, 240
  management concerns, 802.11 networks, 240
  wireless authentication, 236–239
  wireless device portability, 240–241
security
  configure settings in Apache Web Servers, 458–459
  DNS, 442
  ensure physical, 242
  Internet host, 442–443
  mobile device. *See* mobile device security
  router. *See* router security
  routing, 442
security, wireless networking and
  about, 228
  IEEE 802.11, media access control frames, 228–332
  man-in-the-middle attacks, 235
  scanning and attacks, 232–234
  wardriving, exploitation of rogue devices and, 234–235

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

security, wireless solutions, countermeasures
about, 241
AP logging functions, 250–251
AP placement, planning, 242–243
auditing tools, 250
best practices, security, 251–252
default hardware/software settings, changing, 243–244
802.11i advanced encryption standard, 246
encryption and authentication, strong, 244
ensure physical security, 242
instituting security policy, 241–242
mobile device security, 252–254
risk and security assessment, 249
Wi-Fi protected access, 246–249
wireless auditing, 249
security analysis, Web server, 474
Security Association (SA), 394, 422, 423
security auditing, 540, 541
defined, 532–533
independent auditing, 534
operational, 533–534
security-aware resolver, 456, 462
security certifications, 547
sites for, 551
Security Certified Network Architect (SCNA), 551
Security Certified Network Professional (SCNP), 551
Security Certified Program, 540, 551
Security Content Automating Protocol (SCAP), 82
security controls, 483
project, 520–521
security event management, 540
centralized data collection, 528–529
change, managing, 531–532
distributed data collection, 529–530
evaluating IDPS signatures, 530–531
monitoring events, 526–528
multiple sensors, managing data from, 528
program, 526
security event management program, 542
security events, 17
SecurityFocus, 273–274
SecurityFocus Mailing Lists, 550 (table)
security incident response team (SIRT), 507–508, 512, 513
security layers, adding, 537–538
security methods, hardening Windows Web server, 457–458
security policy
acceptable use, employee, 496
approval process, 498
best practices for, 494–495
criminal computer offenses, 496–497
defined, 479, 513
defining yours, 493–494
develop from risk analysis, 495–496
incorporating wireless, 241–242
management and, 497–498
network administration and, 498–499
penalties, outline for violations of, 496
reviews of, conducting routine, 510
risk analysis, use to conduct, 499
security team, information to, 498
updating, 510
security policy, developing
acceptable use, 501–502
extranets and third-party access, 502
Internet/e-mail, secure use of, 504
network security policy, 504–505
physical and facility security, 506
remote access and wireless connections, 503–504
server security, 506
steps to create a, 500–501
user accounts, password protection, logical access control, 502–503
security policy, incident handling procedures
about, 507
escalation procedures, 508–509
responding to security incidents, 509
response team, assembling, 507–508
worst-case scenarios, 509
security policy analysis, conducting, 518
security policy life cycle
about, 477
needs assessment, 477
performance monitoring, 478–479
system design, 478
system implementation, 478
security professionals, Web sites for, 549–550
security testers, rules of, 34
security trends
certifications, 540
mailing lists and newsgroups, 539–540
Web sites for, 539
security workstation, 350–351, 373
selective acknowledgments, 107
semicolons, SQL commands and, 447
semidirectional antenna, 210
sensor(s)
defined, 291
manage data from multiple, 528
network intrusion and, 270–271
sensor locations, combining IDPS, 283
September 11, 2001, 4
SEQ/ACK analysis, TCP elements in packets, 88
sequence control
data frame fields, 232
management frame field, 229
sequence number, TCP elements in packets, 88
sequence number (BE), ICMP header element, 87
sequence number (LE), ICMP header element, 87
Serpent block algorithms, 181
server farm, 350, 373
servers, attacks against web, 443
server security, 506
service, filtering by, 323–325
service network, 15
service set identifier (SSID), 212, 219
session hijacking, 233
session sniping sensor, 278
SHA. *See* Secure Hash Algorithm
shared configuration, 457

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

shared key authentication, 237, 238 (figure), 257
side channel attacks, 179–180, 401, 423
signature(s), 16, 22
defined, 107
evaluating IDPS, 530–531
security policy and, 504
signature analysis, 107
abnormal network traffic signatures, 95–96
about, 83
analyze packet attacks, 85–88
back door, random scan, 97
bad header information, 83–84
defined, 83
Nmapper, 99–100, 117–118
packet attacks, 85
ping sweeps, 96
port scanning, 97
suspicious payload data
Trojan scans, 98–99
Web signatures, 93–95
signature-based IDPS, 268
signature code, writing IDPS, 272
signature-detection IPDS, 273–274
signature detection systems, 267–269, 291
signature files, 13, 22
Simple Mail Transfer Protocol (SMTP), 37, 438
Simple Network Management Protocol (SNMP), 37, 49, 146
AP logging functions and, 250–251
trap logging, 133
single DMZ/two firewalls, 353
single entry point configuration, 409–410, 423
single-packet attacks, 107
single-session attacks, 270
single shared key, 237
site-local unicast address, IPv6 and, 62
site survey, 242–243, 257
planning wireless network, 262
site-to-site VPN, 388, 423
sky wave propagation, 208
smartphones, mobile security, 252–253. *See also* mobile device security
sniffing, 232, 257
wireless, 250
SNMP trap logging, 132, 146
Snort, IDPS, 284–285
create rule, 303–304
exploring packet capture function, 296–298
IDPS function, exploring, 300–303
installing, 295–296
login functions, 298–300
social engineering, 542
defined, 533
threats to network security, 6, 7 (table). *See also* phishing
socket, 8, 22, 330
software
antivirus, 13
assets, 481
change/add, 535
commercial firewalls, 309–310
firewall, 14, 341, 309
freeware proxy servers, 357–358
hardware, changing default settings, 243–244
VPNs, 389
Software Blade, 14
source address, 32-bit IP address, 414
source code, errors in, buffer attacks and, 443–445
source IP address
IP header and, 45
IPv6 core protocol, 57
source port, TCP elements in packets, 88
spam e-mail, 442
spanning port, 277
spear phishing, attacks, 451, 462
split brain DNS architecture, 455, 462
split DNS architecture, 455, 462
split tunneling, 423
spoofing
DNS, 442
MAC address filtering and, 240
spread spectrum, 203–206, 219
SQL injection attacks, 445, 460
defense against, 448–449
query string attacks, 448
Web form attacks, 445–448
Squid for Linux, 358
SSH. *See* Secure Shell
SSL. *See* Secure Sockets Layer
stack smashing attacks, 444
standard IP ACLs, 130–131, 146
standards. *See* wireless networking standards *and entries for specific standards*
star configuration, 407–408, 423
star topography, VPNs and, 407–408
stateful autoconfiguration, 55, 68
stateful packet filtering, 313–314, 330
stateful protocol analysis, 267–269, 269–270, 291
state information, 269, 291
stateless packet filtering, 312–213, 330
state table, 313, 330
static routes, 122–123, 146
configuring, 153–154
storage, manage, 539
stream cipher, 185
stream index, TCP elements in packets, 88
strobe scan, 98–99
Structured Query Language (SQL), attacks and, 445
stub network, 122–123, 146
stub router, 122–123, 146
Sub-7 2.1, Trojan, 98
subnet mask, 37–38, 68
subnetting, 39–42
subpoena, 497, 513
substitution box functions, 160, 185
supernetting, 125–127, 146
survivability, 484, 513
Survivable Network Analysis (SNA), method, 511, 514
steps in, 484–485
Survivable Network Analysis (SNA), risk analysis method, 484
suspicious events, identify
about, 100–101
advanced attacks, 104
falsified IP address, 101
fragmentation abuses, 102–104
illegal flags, 102–103
packet header discrepancies, 101
remote procedure call attacks, 105
TCP/IP packet options, 103
Symantec Security Response, 549 (table)

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

symmetric algorithms, 163, 185
using in Word 2010, 188–189
symmetric block cipher, 160
synchronization flag (SYN), 89
SYN flood, 7 (table), 308
Synlog, 323
SYN scans, 99
SysAdmin, Audit, Network, Security (SANS) Institute, 550 (table)
syslog server, 134, 146
system design, 478
system development life cycle, 477 (figure)
system events, 17
system implementation, 478
system logging, 133

T

TACACS+, 136
target-to-console ratio, 535, 542
TCP headers, Internet Protocol Version 4 (IPv4), 47–48
TCP life cycle, 51–53
TCP three-way handshake, 51–53
team meetings, risk assessment and, 486
Telnet, 37
Telnet firewall rule, 338–339
Temporal Key Integrity Protocol (TKIP), 246–249, 257
Teredo, 329
Terminal Access Controller Access Control System Plus (TACACS+), 133, 146, 236
terrorists, threats to network security, 4
thermal imaging attack, side channel, 180, 401
third party access, security policy and, 502
threat analysis, conducting, 520–521
Threat and Risk Assessment (TRA), 485–486, 511, 514
threat management, installing Gateway 2010, 375–378
Threat Management Gateway (TMG). *See* Microsoft Forefront Threat Management Gateway (TMG)
threats
defined, 479, 514
probability of, 485 (figure)
types of, 481
three-pronged firewall, 347, 373
thresholds, IPDS and, 272
ticket-granting server (TGS), 405–406, 423
ticket-granting ticket (TGT), 405
tier system, of Internet, 439
tiger teams, 532–533, 542
Time to Live (TTL), IP headers and, 45
timing attacks, side channel, 179, 401
Tinkerbell monitoring program, 533, 542
topology, 423
Total Length, IP header, 45
traffic, 17
bottleneck, 276
class, IPv6 core protocol, 56
traffic rate monitoring, stateful protocol analysis, 269
traffic signature
about, 88
FTP signatures, 91–93
normal network traffic and, 89–91
Transmission Control Protocol/Internet Protocol (TCP/IP), 68
address classes, 38
addressing, 37–38
Classless Interdomain Routing (CIDR), 42–43
OSI model and, 36–37
private address ranges, 38–39
subnetting, 39–42
unicasting, multicasting, and broadcasting, 43
variable length subnet masking (VLSM), 42
transmission power, 199
transmission ranges, radio frequency, 196
transmitter address, control frame fields, 231
Transmitter Power Output (TPO), 200
Transport Layer Security (TLS), 164, 185, 403–404, 423
cryptography standards, 175
Triple Data Encryption Standard (3DES) encryption, 400
Triple DES (3DES), 173–174, 185
Trivial File Transfer Protocol (TFTP), 49
Trojan.Asprox, Trojan, 98
Trojan.Mitglieder.h, Trojan, 98
Trojan programs, 3, 5-7 (table), 22, 442–443
Trojan scans, 98–99
TrueCrypt, for virtual volume encryption, 189–190
true negatives, 267–269, 291
true positives, 267–269, 291
trust anchor, 456, 462
TShark
capturing pin sequence in, 89–91
capturing FTP traffic with, 116
setting PATH to work with, 110–111
using to capture ICMP packets, 111–113
tunnel
defined, 423
VPNs and, 387
tunneling protocols, 514
security policy and, 504
understanding, 392–394
tunnel method encryption, 423
tunnel mode, 178–179
two-factor authentication, 504, 514
Twofish, 164
type, ICMP header element, 86

U

Ubuntu Linux, installing, 294–295
UDP port 1701, 393
unauthorized traffic signature, 96
undecillion, 54–55
unicast, Internet Protocol Version 6 (IPv6) and, 62, 68
unicasting, TCP/IP and, 43
Unified Threat Management (UTM), 389, 423
unique local IPv6 unicast addresses, IPv6 and, 62
United Way foundation, 4
UNIX Sendmail program, 84
UNIX services, filtering, 323
Unlicensed National Information Infrastructure (U-NII), 198
Unlicensed Personal Communications Systems, 198
updating, security policy, 510
urgent flag (URG), 89
U.S. Constitution, Fourth Amendment, 496–497

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

U.S. Department of Defense, 4
ARPANET, 438
U.S. Department of Energy Joint Cybersecurity Coordination Center (JC3), 360
U.S. Federal Information Processing Standard (FIPS), 246
U.S. Federal Trade Commission, 3
user accounts
disabling, 363
security policy and, 502–503
User Data Protocol (UDP) headers, Internet Protocol Version 4 (IPv4), 48–49, 68
user management, IDPS, 535
utilities, IPv6 and, 63
UUNET/MCI World-com, 438

V

vanilla scan, 98, 107
variable length subnet masking (VLSM), 42, 68, 126, 146
VeriSign, 459
virtual private network (VPN), 9, 22, 423, 504
about, 386
accessing internal network from, 429–431
advantages and disadvantages of, 391
auditing compliance, 419–420
authentication, 404–406
Authentication Header and Encapsulating Security Payload, 396
business needs, evaluating, 389–391
components, 387–388
configure Microsoft Forefront Threat Management Gateway (TMG) for, 427–429
defined, 386–387
domains, 409
encapsulation, 391–392
encryption, 398–399
encryption schemes, 399–404
ensuring client security, 416–417
filtering rules, 435
firewalls, using with, 411–414
hardware, 388–389
hybrid topography, 408–409
logging activity on, 419
mesh topology, 406–407
network security and, 16–17
packet filtering rules, adjust for, 414–416
policies, guidelines, 420–421
quarantine, 418–419
remote access, create a domain group for, 426–427
remote access, security and, 10 (figure)
single/multiple entry configurations, 409–411
software, 389
star topography, 407–408
Transport Layer Security (TLS), 403–404
Triple Data Encryption Standard (3DES) encryption, 400
tunneling protocols, understanding, 392–394
virtual terminal, 146
virtual volume encryption, using TrueCrypt, 189–190
virus, 442–443
defined, 22
threats to network security, 5, 7 (table)
scanning, 22
Visual Basic, ActiveX control attacks and, 452
Voice over Internet Protocol (VoIP), 214, 455
Voice over Wi-FI (VoWiFi), 214
voltage standing wave ratio (VSWR), 198, 199, 219–220
VPN
appliance, 388, 423
client, 387, 423
domain, 409, 423
protocols, 387, 424
quarantine, 424
server, 387, 424 *See also* virtual private network
VTY passwords, 138
vulnerability, 514
defined, 479, 482
evaluate, 486
vulnerability status, 83
vulnerable services, 18
detecting with CVE, 83

W

wardriving, 257
defined, 234
exploitation of rogue devices and, 234–235
watt, 199, 220
wave, 194, 199, 220
wavelength, frequency and, 195, 220
wave motion, 208–209
wave propagation, 209, 220
Web access policy, configuring Web, 379–381
Web application attacks, 8 (table)
Web-based attacks, discover, 474
Web form attacks, SQL injection and, 445–448
Web infrastructure DMZ, 455 (figure)
Web page deception, attacks against users, 449–451
Web servers, attack techniques against, 443
Web server security analysis, conducting, 474
Web service rules, packet filtering, 325
Web signatures, analyzing, 93–95
Web site attack techniques
against Web servers, 443
against Web users, 449–453. *See also* attacks, against Web users
buffer overflow attacks, 443–445
SQL injection, 445–449. *See also* SQL injection attacks
Web site(s)
blocking, 383–384
for security, 539
Web user(s), attacks against, 449, 460
WepAttack, 233
whitelists, IDPS and, 272
Wi-Fi Protected Access Version 2 (WPA2), 247–249, 257
Wi-Fi Protected Access (WPA), 164, 246–249, 257
Wi-Fi Protected Setup (WPS), 248, 257
Windows
exploring with IPv6 implementation with Netsh, 154–156
filtering and, 321–323
firewall, advanced settings of, 332–334

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.

Windows Basic Authentication, 457
Windows Digest Authentication, 457
Windows Firewall with Advanced Security window, 315
window size value, TCP elements in packets, 88
Windows Network Monitor, 45
Windows operating systems
installing inSSIDer on, 223–225
IPv6 configuration and, 63
malicious codes, threats to network security, 5
packet filtering and, 14
Windows Server 2008, 360
installing FTP server on, 113–115
Windows Web server, hardening, 456–458 *See also* hardening Windows Web servers
Wired Equivalent Privacy (WEP), 164, 182, 229
default keys, 239
vulnerabilities of, 237–239
wireless auditing, 249
wireless authentication, 236–239
wireless communications
about, 194
analog modulation, 201
digital modulation, 201–203
electromagnetic radiation, 194–195
infrared transmissions, 195
interference radio frequencies, 197
measuring RF signals and, 199–200
radio frequency signal behavior, 197–198
radio frequency transmissions, 196
spread spectrum, 203–206
transmission ranges, 196–197
wireless connections, security policy and, 503–504
wireless device
finding default values for, 260
observe how people handle, 261–262
physical security for, 260
wireless device portability, 240–241
wireless gateways, 211–212
wireless LANS, components
access points, 206–208
antennas, 208–210
NICs, 206
remote wireless bridge, 210–211
wireless gateways, 211–212
WLAN configurations, 212–213
wireless networking, defined, 220
wireless networking security, *See* security, wireless networking
wireless networking standards, 213–215, 236–237
wireless networking utility, installing inSSIDer, 223–225
Wireless Network Management amendment, 215
wireless security solutions and countermeasures. *See* security, wireless solutions, countermeasures
wireless signal interference, researching, 226
wireless sniffers, 250
Wireshark, 46, 250
analyzing packet attacks with, 85–88
examining IPv6 ping packets, 76–77
ICMP packets, use TShark to capture, 111–113
individual ARP and ping packets in, 73–75
installing, 71
setting PATH variable to work with TShark, 110–111
use to capture IPv4 pings, 71–73
WLAN configuration, 212–213
WLAN implementation, secure. *See* secure WLAN implementation
Word 2010
configuring encryption in, 192
using a symmetric algorithm in, 188–189
World Trade Center, 4
World Wide Web, 438
Web server attacks and, 443
worm, threats to network security, 5, 22
worst-case scenarios, 509, 514
WPA2-AES Enterprise Security, 248
WPA2-AES Personal Security, 248
WPA Cracker, 233, 250
W32.Spybot.pen, Trojan, 98

X

X.509, PKI, 172–173, 185
Xmas scans, 100
XOR function, 159–160, 403
XSL attacks, 181, 401

Z

zombie computers, 442
ZoneAlarm Internet Security Suite, 14, 306, 310
firewall, 309
installing, 25–27
viewing log files, 30–33
zone file, 454
zones, of defense, 537–538
zone transfers, 454–455, 462

Copyright 2012 Cengage Learning. All Rights Reserved. May not be copied, scanned, or duplicated, in whole or in part. Due to electronic rights, some third party content may be suppressed from the eBook and/or eChapter(s). Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. Cengage Learning reserves the right to remove additional content at any time if subsequent rights restrictions require it.