

“All-in-One Is All You Need.”

ALL-IN-ONE

SSCP[®]

Systems Security Certified Practitioner

EXAM GUIDE
SECOND EDITION



Electronic content includes:

- Two practice exams
- PDF copy of the book

*Complete coverage of
all seven exam domains
of the SSCP exam*



*Includes updated exam
objectives effective April 15, 2015*



*Ideal as both a study tool and
an on-the-job reference*



*Filled with practice exam
questions and in-depth
explanations*

**Mc
Graw
Hill**
Education

DARRIL GIBSON

SSCP, CISSP[®], CompTIA Security+[™], CASP, MCSE, MCITP, ITIL[®]

www.allitebooks.com

ALL ■ IN ■ ONE

SSCP[®]
Systems Security
Certified Practitioner
EXAM GUIDE

Second Edition

Darril Gibson



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

McGraw-Hill Education is an independent entity from (ISC)²® and is not affiliated with (ISC)² in any manner. This study/training guide and/or material is not sponsored by, endorsed by, or affiliated with (ISC)² in any manner. This publication and digital content may be used in assisting students to prepare for the SSCP[®] exam. Neither (ISC)² nor McGraw-Hill Education warrant that use of this publication and digital content will ensure passing any exam. (ISC)²®, CISSP®, CAP®, ISSAP®, ISSEP®, ISSMP®, SSCP®, CSSLP®, and CBK® are trademarks or registered trademarks of (ISC)² in the United States and certain other countries. All other trademarks are trademarks of their respective owners.

Copyright © 2016 by McGraw-Hill Education. All rights reserved. Except as permitted under the United States 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 the 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.

ISBN: 978-1-25-958306-3

MHID: 1-25-958306-6

The material in this eBook also appears in the print version of this title: ISBN: 978-1-25-958307-0,
MHID: 1-25-958307-4.

eBook conversion by codeMantra
Version 1.0

All trademarks are trademarks of their respective owners. Rather than put a trademark symbol after every occurrence of a trademarked name, we use names in an editorial fashion only, and to the benefit of the trademark owner, with no intention of infringement of the trademark. Where such designations appear in this book, they have been printed with initial caps.

McGraw-Hill Education eBooks 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 page at www.mhprofessional.com.

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.

TERMS OF USE

This is a copyrighted work and McGraw-Hill Education and its licensors reserve all rights in and to the work. Use of this work is subject to these terms. Except as permitted under the Copyright Act of 1976 and the right to store and retrieve one copy of the work, you may not decompile, disassemble, reverse engineer, reproduce, modify, create derivative works based upon, transmit, distribute, disseminate, sell, publish or sublicense the work or any part of it without McGraw-Hill Education's prior consent. You may use the work for your own noncommercial and personal use; any other use of the work is strictly prohibited. Your right to use the work may be terminated if you fail to comply with these terms.

THE WORK IS PROVIDED "AS IS." MCGRAW-HILL EDUCATION AND ITS LICENSORS MAKE NO GUARANTEES OR WARRANTIES AS TO THE ACCURACY, ADEQUACY OR COMPLETENESS OF OR RESULTS TO BE OBTAINED FROM USING THE WORK, INCLUDING ANY INFORMATION THAT CAN BE ACCESSED THROUGH THE WORK VIA HYPERLINK OR OTHERWISE, AND EXPRESSLY DISCLAIM ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. McGraw-Hill Education and its licensors do not warrant or guarantee that the functions contained in the work will meet your requirements or that its operation will be uninterrupted or error free. Neither McGraw-Hill Education nor its licensors shall be liable to you or anyone else for any inaccuracy, error or omission, regardless of cause, in the work or for any damages resulting therefrom. McGraw-Hill Education has no responsibility for the content of any information accessed through the work. Under no circumstances shall McGraw-Hill Education and/or its licensors be liable for any indirect, incidental, special, punitive, consequential or similar damages that result from the use of or inability to use the work, even if any of them has been advised of the possibility of such damages. This limitation of liability shall apply to any claim or cause whatsoever whether such claim or cause arises in contract, tort or otherwise.

To my wife Nimfa—
Thanks for sharing your life with me for the past 23 years and
letting me share mine with you.

ABOUT THE AUTHOR

Darril Gibson is the CEO of YCDA, LLC (short for You Can Do Anything) and he has authored or coauthored more than 35 books. Darril regularly writes, consults, and teaches on a wide variety of technical and security topics and holds several certifications, including (ISC)² SSCP and CISSP; CompTIA Security+ and CASP; Microsoft MCSE and MCITP; and ITIL Foundations. In response to repeated requests, Darril created the <http://gcgapremium.com/> site where he provides additional study materials for several certification exams. He regularly posts blog articles (<http://blogs.getcertifiedgetahead.com/>) about certification topics and uses that site to help people stay abreast of changes in certification exams. You can contact him through either of these sites. Darril lives in Virginia Beach with his wife and two dogs. Whenever possible, they escape to a small cabin in the country on over 20 acres of land that continues to provide them with peace, tranquility, and balance.

About the Technical Editor

Josh More has more than 15 years of experience in security, IT, development, and system and network administration. Currently, he runs Eyra Security, a security and business improvement consulting firm based in Minneapolis, MN. Josh holds several security and technical certifications and has served in a leadership position on several security-focused groups. He writes a blog on security at www.starmind.org, often taking a unique approach to solving security problems by applying lessons from other disciplines like agile development, lean manufacturing, psychology, economics and complexity science. He has also written several books on IT, information security, and career management.

CONTENTS AT A GLANCE

Chapter 1	Security Fundamentals	1
Chapter 2	Access Controls	27
Chapter 3	Basic Networking and Communications	67
Chapter 4	Advanced Networking and Communications	115
Chapter 5	Attacks	157
Chapter 6	Malicious Code and Activity	207
Chapter 7	Risk, Response, and Recovery	243
Chapter 8	Monitoring and Analysis	275
Chapter 9	Controls and Countermeasures	303
Chapter 10	Auditing	343
Chapter 11	Security Operations	371
Chapter 12	Security Administration and Planning	407
Chapter 13	Legal Issues	439
Chapter 14	Cryptography	465
Appendix	About the Download	511
	Glossary	513
	Index	539

This page intentionally left blank

CONTENTS

Acknowledgments	xviii
Introduction	xix
Chapter 1 Security Fundamentals	1
Reviewing the Requirements for SSCP	1
Registering for the Exam	1
Have One Year of Experience	3
Passing the Exam	4
Maintaining Your SSCP Certification	7
Understanding Basic Security Concepts	8
Confidentiality	9
Integrity	10
Availability	12
Exploring Fundamentals of Security	13
Least Privilege	13
Separation of Duties	14
Privacy	15
Defense in Depth	15
Nonrepudiation	16
AAAs of Security	17
Accountability	18
Due Diligence	19
Due Care	19
Chapter Review	20
Questions	21
Answers	23
Chapter 2 Access Controls	27
Comparing Identification, Authentication, and Authorization	27
Exploring Authentication	28
Three Factors of Authentication	29
Multifactor Authentication	37
Reviewing Identification	38
Single Sign-on Authentication	38
Centralized vs. Decentralized Authentication	42
Offline Authentication	43
Device Authentication	43

Implementing Access Controls	44
Comparing Subjects and Objects	44
Logical Access Controls	47
Comparing Access Control Models	47
Discretionary Access Control	47
Non-Discretionary Access Control	49
Access Control Matrix vs. Capability Table	55
Participating in the Identity-Management Life Cycle	55
Identity Proofing	56
Provisioning and Authorization	56
Maintenance and Entitlement	57
De-provisioning	58
Participating in Physical Security Operations	58
Chapter Review	59
Questions	61
Answers	64
Chapter 3 Basic Networking and Communications	67
The OSI Model	67
The Physical Layer (Layer 1)	68
The Data Link Layer (Layer 2)	69
The Network Layer (Layer 3)	70
The Transport Layer (Layer 4)	70
The Session Layer (Layer 5)	71
The Presentation Layer (Layer 6)	72
The Application Layer (Layer 7)	72
Comparing the OSI and TCP/IP Models	72
Network Topologies	73
Ethernet	73
Bus	75
Star	76
Tree	77
Token Ring	77
Mesh	78
Reviewing Basic Protocols and Ports	79
Comparing IPv4 and IPv6	79
Dynamic Host Configuration Protocol	80
Address Resolution Protocol	81
Network Discovery Protocol	82
Domain Name System	82
Internet Control Message Protocol	83
Internet Group Message Protocol	83
Simple Network Management Protocol	84
File Transfer Protocol	84
Telnet	85

Secure Shell	85
HyperText Transfer Protocol and HyperText Transfer Protocol Secure	86
Transport Layer Security and Secure Sockets Layer	86
Network File System	87
Routing Protocols	87
E-mail Protocols	87
Tunneling Protocols	88
Internet Protocol Security	88
Mapping Well-Known Ports to Protocols	89
Comparing Ports and Protocol Numbers	91
Comparing Internetwork Trust Architectures	91
Comparing Public and Private IP Addresses	93
Using NAT	94
Comparing Trust Relationships	96
Exploring Wireless Technologies	97
Securing Data Transmissions	99
Wireless Device Administrator Password	101
Wireless Service Set Identifier	102
MAC Filtering	103
Bluetooth	104
GSM	104
3G, LTE, and 4G	104
WiMAX	105
Radio Frequency Identification	105
NFC	105
Protecting Mobile Devices	106
Chapter Review	107
Questions	109
Answers	112
Chapter 4 Advanced Networking and Communications	115
Managing LAN-Based Security	115
Comparing Switches and Routers	115
Segmentation	117
Secure Device Management	120
Understanding Telecommunications	120
Internet Connections	120
VoIP	122
Securing Phones	122
Converged Communications	123
Using Proxy Servers	123
Understanding Firewalls	125
Packet-Filtering Firewall	125
Stateful Inspection Firewall	127

	Application Firewall	127
	Next-Generation Firewall	128
	Defense Diversity	128
	Comparing Network-based and Host-based Firewalls	129
	Exploring Remote Access Solutions	130
	Risks and Vulnerabilities	131
	Tunneling Protocols	131
	Authentication	134
	Traffic Shaping	138
	Access and Admission Control	138
	Exploring Virtual Environments	140
	Virtualization Terminology	140
	Shared Storage	141
	Virtual Appliances	141
	Continuity and Resilience	142
	Separation of Data Plane and Control Plane	142
	Software-defined Networking	143
	Attacks and Countermeasures	143
	Understanding Cloud Computing	144
	Cloud Operation Models	145
	Storage	146
	Privacy	147
	Data Control and Third-party Outsourcing	147
	Compliance	148
	Chapter Review	148
	Questions	150
	Answers	153
Chapter 5	Attacks	157
	Comparing Attackers	157
	Hackers and Crackers	158
	White Hats, Black Hats, and Grey Hats	158
	Advanced Persistent Threats	159
	Insider Attacks	160
	Script Kiddies	161
	Phreaks	162
	Accidental Threats	162
	Exploring Attack Types and Countermeasures	163
	Basic Countermeasures	163
	Spoofing	163
	DoS	164
	DDoS	165
	Botnets and Zombies	165
	Sniffing Attack	167
	Ping Sweep	170

Port Scan	170
Salami Attack	171
Man-in-the-Middle	171
Session Hijacking	172
Replay	173
Smurf and Fraggle Attacks	173
Software Security as a Countermeasure	174
Buffer Overflow Attacks	177
Injection Attacks	178
Cross-Site Scripting	180
Cross-Site Request Forgery	180
Password Attacks	182
Spam	185
Phishing Attacks	185
Phishing and Drive-by Downloads	187
Spear Phishing and Whaling	187
Vishing	188
Smishing	188
Zero Day Exploits	188
Covert Channel	190
Wireless Attacks and Countermeasures	190
Understanding Social Engineering	193
Tailgating	194
Impersonation	195
Dumpster Diving	195
Shoulder Surfing	195
Pharming	195
Social Networking Attacks	196
User Awareness as a Countermeasure	196
Chapter Review	197
Questions	199
Answers	202
Chapter 6 Malicious Code and Activity	207
Identifying Malicious Code	207
Virus	207
Worm	210
Trojan Horse	211
Scareware	211
Ransomware	213
Keylogger	215
Logic Bomb	215
Rootkits	215
Mobile Code	216
Backdoors and Trapdoors	217

	RATs	218
	Spyware	218
	Malware Hoaxes	218
	Analyzing the Stages of Regain	219
	Understanding Malware Delivery Methods	221
	Delivering Malware via Drive-by Downloads	221
	Delivering Malware via Malvertising	222
	Delivering Malware via E-mail	223
	Delivering Malware via USB Drives	223
	Implementing Malicious Code Countermeasures	223
	Antivirus Software	224
	Keeping AV Signatures Up to Date	228
	Spam Filters	229
	Content-Filtering Appliances	229
	Keeping Operating Systems Up to Date	231
	Scanners	231
	Beware of Shortened Links	231
	Sandboxing	232
	Least Privilege	232
	Software Security	233
	Application Whitelisting and Blacklisting	234
	Participating in Security Awareness and Training	234
	Common Vulnerabilities and Exposures	235
	Chapter Review	235
	Questions	236
	Answers	240
Chapter 7	Risk, Response, and Recovery	243
	Defining Risk	243
	Identifying Threat Sources	245
	Identifying Threat Events	246
	Understanding Vulnerabilities	248
	Understanding Impact	249
	Managing Risk	250
	Residual Risk	251
	Identifying Assets	252
	Risk Visibility and Reporting	253
	Risk Register	253
	Performing Risk Assessments	254
	Quantitative Analysis	254
	Qualitative Analysis	256
	Risk Assessment Steps	258
	Address Findings	262
	Responding to Incidents	262
	Preparation	263

	Detection and Analysis	265
	Containment, Eradication, and Recovery	266
	Post-incident Activity	267
	Chapter Review	267
	Questions	268
	Answers	272
Chapter 8	Monitoring and Analysis.....	275
	Operating and Maintaining Monitoring Systems	275
	Intrusion Detection Systems	275
	IDS Alerts	276
	Network-based Intrusion Detection Systems	277
	Host-based Intrusion Detection Systems	278
	Intrusion Prevention Systems	279
	Detection Methods	282
	Wireless Intrusion Detection and Prevention Systems	283
	Analyzing Results	283
	Detection Systems and Logs	284
	Detecting Unauthorized Changes	284
	Using Security Information and Event Management Tools	286
	Performing Security Assessment Activities	287
	Vulnerability Assessments	287
	Penetration Tests	294
	Chapter Review	296
	Questions	297
	Answers	300
Chapter 9	Controls and Countermeasures	303
	Using Controls, Safeguards, and Countermeasures	303
	Performing a Cost-Benefit Analysis	304
	Security Controls Life Cycle	305
	Understanding Control Goals	307
	Preventive	307
	Detective	308
	Corrective	309
	Other Controls	309
	Comparing the Classes of Controls	311
	Management/Administrative Security Controls	311
	Technical Security Controls	312
	Operational Security Controls	312
	Physical Security Controls	313
	Combining Control Goals and Classes	313
	Exploring Some Basic Controls	314
	Hardening Systems	314
	Policies, Standards, Procedures, and Guidelines	315

Response Plans	317
Change Control and Configuration Management	317
Testing Patches, Fixes, and Updates	318
Endpoint Device Security	320
User Awareness and Training Programs	325
Understanding Fault Tolerance	325
Fault Tolerance for Disks	325
Failover Clusters	329
Redundant Connections	330
Understanding Backups	331
Full Backups	332
Full/Incremental Backup Strategy	332
Full/Differential Backup Strategy	333
Chapter Review	334
Questions	335
Answers	339
Chapter 10 Auditing	343
Understanding Auditing and Accountability	343
Holding Users Accountable with Audit Logs	344
Auditing with Logs	345
Clipping Levels	346
Understanding Audit Trails	348
Exploring Audit Logs	348
Operating System Logs	348
Storing Logs on Remote Systems	349
*Nix Logs	350
Proxy Server Logs	351
Firewall Logs	352
Reviewing Logs	352
Managing Audit Logs	353
Performing Security Audits	354
Auditing Passwords	355
Auditing Security Policies	355
ISACA	356
Exploring PCI DSS Requirements	356
Auditing Physical Access Controls	358
Understanding Configuration Management	358
Using Imaging for Configuration Management	359
Using Group Policy for Configuration Management	360
Understanding Change Management	361
Chapter Review	363
Questions	363
Answers	367

Chapter 11	Security Operations	371
	Handling Data	371
	Classifying Data	371
	Marking and Labeling Data	374
	Roles and Responsibilities	374
	Protecting Data from Cradle to Grave	375
	Data at Rest and Data in Motion	375
	Data Management Policies	376
	Understanding Databases	382
	Data Inference	386
	Data Diddling	386
	Securing Big Data	387
	Regulatory Requirements	387
	Training	390
	Managing Assets	390
	Hardware	390
	Software	391
	Data	391
	Certification and Accreditation	392
	Certification, Accreditation, and Security Assessments	392
	Common Criteria	393
	Using a Risk Management Framework	394
	Understanding Security Within the System Development Life Cycle	395
	Chapter Review	398
	Questions	399
	Answers	402
Chapter 12	Security Administration and Planning	407
	Understanding Security Policies	407
	Security Policy Characteristics	408
	Enforcing Security Policies	412
	Value of a Security Policy	412
	Security Policies Becoming More Common	413
	Understanding Code of Ethics	414
	Policy Awareness	415
	Updating Security Policies	416
	Understanding BCPs and DRPs	417
	Business Impact Analysis	419
	Disaster Recovery Plan	422
	Emergency Response Plans and Procedures	423
	Comparing a BCP and a DRP	423
	Restoration Planning	424
	Testing and Drills	424
	Alternative Locations	425

Identifying Security Organizations	428
NIST	428
US-CERT	429
SANS Institute	430
CERT Division	430
Chapter Review	430
Questions	431
Answers	435
Chapter 13 Legal Issues	439
Exploring Computer Forensics	439
Participating in Incident Handling	439
First Responders and Preserving the Scene	442
Three Phases of a Computer Forensics Investigation	443
Forensic Evidence Guidelines and Principles	447
Comparing Computer Abuse and Computer Crime	448
Understanding Fraud and Embezzlement Crime	450
Mandatory Vacations	450
Job Rotation	451
Understanding Privacy Issues	452
European Directives	454
California Supreme Court Rules That ZIP Codes Are PII	455
Connecticut's Public Act No. 08-167	455
Children's Online Privacy Protection Act	456
California Online Privacy Protection Act of 2003	456
Chapter Review	457
Questions	458
Answers	461
Chapter 14 Cryptography	465
Understanding Basic Cryptography Concepts	465
Cryptography Terminology	466
Data Sensitivity	467
Regulatory Requirements	468
Participating in Security Awareness and Training	469
Enforcing Integrity with Hashing	469
Hashing Algorithms Provide One-Way Encryption	469
Hashing Algorithms	470
Verifying a Hash	472
Salting Passwords	473
Exploring Symmetric Encryption	474
ROT13	475
Composing and Rotating Keys	475
Comparing Block and Stream Ciphers	476

Advanced Encryption Standard	477
Other Symmetric Encryption Algorithms	477
Exploring Asymmetric Encryption	479
RSA	481
Transport Layer Security	481
Secure Sockets Layer	483
Diffie-Hellman	483
Elliptic Curve Cryptography	484
Secure Shell	484
Protecting E-mail with S/MIME	484
Pretty Good Privacy (PGP)	490
Other Encryption Schemes	490
Steganography	490
IPsec	492
Public Key Infrastructure	492
Certificates	492
Certificate Authority	495
Key Escrow	499
Alternative Certificate Trusts	500
Comparing Cryptanalysis Attacks	501
Managing Cryptographic Keys	501
Known-Plaintext Attack	502
Ciphertext-Only Attack	502
Chapter Review	502
Questions	504
Answers	507
Appendix About the Download	511
System Requirements	511
Downloading Total Tester Premium Practice Exam Software	511
Total Tester Premium Practice Exam Software	511
Installing and Running Total Tester	512
Technical Support	512
Total Seminars Technical Support	512
McGraw-Hill Education Content Support	512
Glossary	513
Index	539

ACKNOWLEDGMENTS

Books are never done alone but instead are a result of collaboration among many people. I'm very grateful for the hard work done by several people on this project, including Tim Green, who had the faith in me to write the book; Meghan Manfre and Amy Stonebraker, who helped keep the project on track; and Josh More, the technical editor who provided some thoughtful feedback. I'm especially grateful to the copy editor, Bill McManus, who expended a great deal of time and energy on this book. Not only did he do standard copyediting, he also helped identify and correct several technical issues within the content. Thanks again, Bill. And, of course, I'm very grateful for the support my wife gave me as I worked on this update, giving me the time and space I needed to finish it in a timely manner.

INTRODUCTION

The importance of information technology (IT) security increases every day. The news is filled with reports of data breaches where customer data is lost and companies are forced to scramble with crisis management. Rarely a day goes by when users don't see phishing e-mails in their inbox, with each phishing e-mail trying to trick them into clicking an unsafe link or giving up valuable information. Individual botnets commonly control tens of thousands of computers, all ready at a moment's notice to launch attacks.

All of these risks can cause substantial losses for an organization. Instead of waiting for an incident and then responding, organizations are realizing they must be more proactive with IT security. The alternative is to do nothing and then watch profits slip away.

More IT security jobs have become available to fill this void, but employers often find it difficult to locate professionals with the right mix of knowledge and skills, including security knowledge and skills. Hiring managers typically want to have some indication that the person they're hiring for an IT job has at least some security knowledge, and that's where certifications fit in. Individuals with a security certification such as the Systems Security Certified Practitioner (SSCP) will often get the job interview, while individuals without a security certification may be overlooked.

This book can help you learn the material to prepare for the SSCP exam. Just as importantly, it can help you build your IT security knowledge so that you can provide real and lasting assistance to organizations that are seeking to improve their security posture.

Exam Objective Map

The SSCP exam is composed of questions from seven domains:

- Access Controls
- Security Operations and Administration
- Risk Identification, Monitoring, and Analysis
- Incident Response and Recovery
- Cryptography
- Network and Communications Security
- Systems and Application Security

The following table maps the domain topics for the SSCP exam to the chapter and section where each is covered.

Domain	Chapter Number: Section(s)
1) Access Controls	
Implement Authentication Mechanisms	Chapter 2: Exploring Authentication
Operate Internetwork Trust Architectures	Chapter 3: Comparing Internetwork Trust Architectures
Participate in the Identity-Management Life Cycle	Chapter 2: Participating in the Identity-Management Life Cycle
Implement Access Controls	Chapter 2: Comparing Access Control Models
2) Security Operations and Administration	
Understand and Comply with Codes of Ethics	Chapter 1: Reviewing the Requirements for SSCP Chapter 12: Understanding Security Policies
Understand Security Concepts	Chapter 1: Understanding Basic Security Concepts; Exploring Fundamentals of Security
Document and Operate Security Controls	Chapter 9: Understanding Control Goals
Participate in Asset Management	Chapter 11: Managing Assets
Implement and Assess Compliance with Controls	Chapter 9: Comparing the Classes of Controls
Participate in Change Management	Chapter 9: Exploring Some Basic Controls Chapter 10: Understanding Configuration Management; Understanding Change Management
Participate in Security Awareness and Training	Chapter 5: Exploring Attack Types and Countermeasures; Understanding Social Engineering Chapter 6: Implementing Malicious Code Countermeasures Chapter 9: Exploring Some Basic Controls Chapter 11: Handling Data Chapter 12: Understanding Security Policies Chapter 14: Understanding Basic Cryptography Concepts
Participate in Physical Security Operations	Chapter 2: Participating in Physical Security Operations Chapter 9: Comparing the Classes of Controls Chapter 10: Auditing Physical Access Controls
3) Risk Identification, Monitoring, and Analysis	
Understand the Risk Management Process	Chapter 7: Defining Risk; Managing Risk; Performing Risk Assessments
Perform Security Assessment Activities	Chapter 7: Performing Risk Assessments Chapter 8: Operating and Maintaining Monitoring Systems; Performing Security Assessment Activities
Operate and Maintain Monitoring Systems	Chapter 8: Operating and Maintaining Monitoring Systems; Using Security Information and Event Management Tools
Analyze Monitoring Results	Chapter 8: Operating and Maintaining Monitoring Systems

Domain	Chapter Number: Section(s)
4) Incident Response and Recovery	
Participate in Incident Handling	Chapter 7: Responding to Incidents Chapter 13: Exploring Computer Forensics
Understand and Support Forensic Investigations	Chapter 13: Exploring Computer Forensics
Understand and Support Business Continuity Plan (BCP) and Disaster Recovery Plan (DRP)	Chapter 9: Understanding Fault Tolerance; Understanding Backups Chapter 12: Understanding BCPs and DRPs
5) Cryptography	
Understand and Apply Fundamental Concepts of Cryptography	Chapter 14: Understanding Basic Cryptography Concepts; Enforcing Integrity with Hashing; Exploring Symmetric Encryption; Exploring Asymmetric Encryption
Understand Requirements for Cryptography	Chapter 11: Handling Data Chapter 14: Understanding Basic Cryptography Concepts
Understand and Support Secure Protocols	Chapter 3: Reviewing Basic Protocols and Ports Chapter 4: Exploring Remote Access Solutions Chapter 14: Exploring Symmetric Encryption; Exploring Asymmetric Encryption; Other Encryption Schemes
Operate and Implement Cryptographic Systems	Chapter 14: Exploring Symmetric Encryption; Exploring Asymmetric Encryption; Other Encryption Schemes; Public Key Infrastructure
6) Network and Communications Security	
Understand Security Issues Related to Networks	Chapter 3: The OSI Model; Comparing the OSI and TCP/IP Models; Network Topologies; Reviewing Basic Protocols and Ports Chapter 4: Access and Admission Control
Protect Telecommunications Technologies	Chapter 4: Understanding Telecommunications Chapter 5: Exploring Attack Types and Countermeasures
Control Network Access	Chapter 4: Access and Admission Control
Manage LAN-Based Security	Chapter 4: Managing LAN-Based Security
Operate and Configure Network-Based Security Devices	Chapter 4: Using Proxy Servers; Understanding Firewalls; Exploring Remote Access Solutions Chapter 8: Operating and Maintaining Monitoring Systems
Implement and Operate Wireless Technologies	Chapter 3: Exploring Wireless Technologies Chapter 5: Exploring Attack Types and Countermeasures Chapter 8: Operating and Maintaining Monitoring Systems

Domain	Chapter Number: Section(s)
7) Systems and Application Security	
Identify and Analyze Malicious Code and Activity	Chapter 5: Comparing Attackers; Exploring Attack Types and Countermeasures; Understanding Social Engineering Chapter 6: Identifying Malicious Code; Identifying Malware Delivery Methods; Implementing Malicious Code Countermeasures
Implement and Operate Endpoint Device Security	Chapter 4: Understanding Firewalls Chapter 6: Implementing Malicious Code Countermeasures Chapter 8: Operating and Maintaining Monitoring Systems Chapter 9: Endpoint Device Security
Operate and Configure Cloud Security	Chapter 4: Understanding Cloud Computing
Secure Big Data Systems	Chapter 11: Handling Data
Operate and Secure Virtual Environments	Chapter 4: Exploring Virtual Environments

Security Fundamentals

In this chapter, you will learn about

- Requirements to earn the (ISC)² Systems Security Certified Practitioner (SSCP) certification
- Primary goals of security related to confidentiality, integrity, and availability (CIA)
- Fundamental security terminology

Reviewing the Requirements for SSCP

The Systems Security Certified Practitioner (SSCP) certification is one of the certifications sponsored by the International Information Systems Security Certification Consortium, Inc., more commonly known as (ISC)².

There are several requirements that you must complete to earn the certification:

- Have at least one year of experience in one or more of the (ISC)² SSCP domains
- Legally commit to abide by the (ISC)² Code of Ethics
- Answer four questions regarding criminal history and related background
- Pass the exam



EXAM TIP Earning the certification is more than just passing the exam. You must also have one year of experience and commit to the Code of Ethics.

If you've earned certifications from other vendors (such as CompTIA), you'll find that the (ISC)² process is different. It is often confusing to people the first time they earn an (ISC)² certification. The following sections explain the process of registering for the exam, some key information about the exam, the requirement to submit paperwork to validate your experience, and requirements to maintain the SSCP certification.

Registering for the Exam

The exams are computer based and administered at Pearson VUE test centers. You can register for exams through the Pearson VUE website (www.pearsonvue.com/isc2/). If you register online, you'll need to have or create a Pearson VUE account. The Pearson VUE website also includes a search feature so that you can locate a test center near you.



NOTE The SSCP exam was previously paper based and proctored in large conference rooms. While (ISC)² has moved to Computer-Based Testing (CBT), it does occasionally authorize paper-based exams on a limited basis, such as after an official training seminar. The primary portal for all SSCP information (www.isc2.org/sscp) includes up-to-date information on the exam, including current exam prices.

Registering for the exam includes three steps:

1. Submit the exam fee.
2. Legally commit to abide by the (ISC)² Code of Ethics.
3. Answer four questions on criminal history and related background.

Submit the Exam Fee

You submit the exam fee through Pearson VUE, the testing provider. This page (www.vue.com/isc2/) includes links you can use to create a Pearson VUE account, find a testing center close to you, and register for the exam. Some organizations purchase vouchers in bulk and give them to their employees. If you have a voucher, you can use it instead of submitting a fee.

Committing to Abide by the (ISC)² Code of Ethics

The Code of Ethics includes a preamble and four canons describing (ISC)²'s ethical expectations of its certified practitioners. Candidates must commit to and abide by the Code of Ethics to earn and keep the SSCP certification. Members who violate any provision of the Code of Ethics may have their certification revoked based on recommendations from a peer review panel.

The following sections quote the preamble and canons exactly as they appear on the (ISC)² page (<https://www.isc2.org/ethics>).

Code of Ethics Preamble The preamble consists of two points:

- The safety and welfare of society and the common good, duty to our principals, and to each other, requires that we adhere, and be seen to adhere, to the highest ethical standards of behavior.
- Therefore, strict adherence to this Code is a condition of certification.

Code of Ethics Canons The four canons are as follows:

- Protect society, the common good, necessary public trust and confidence, and the infrastructure.
- Act honorably, honestly, justly, responsibly, and legally.
- Provide diligent and competent service to principals.
- Advance and protect the profession.

If a situation arises resulting in a conflict between the canons, the conflict should be resolved in the order in which the canons are listed. In other words, the first canon is more important than the second one, and so forth.



EXAM TIP The (ISC)² Code of Ethics is included in the Security Operations and Administration domain and you can expect to be tested on it.

Answering Questions Related to Criminal History and Background

The registration process requires you to answer several questions related to your history and background. These questions ask about the following topics:

- Felony convictions
- Involvement with hackers or hacking
- Revocation of any licenses or certifications
- Any use of aliases or pseudonyms

Answering yes to any of these questions doesn't necessarily disqualify you. However, it's best to resolve a potential problem before taking the exam. You can e-mail (ISC)² at legal@isc2.org to discuss your situation. If appropriate, (ISC)² will declare you eligible to take the exam and pursue the certification.

Have One Year of Experience

To earn the SSCP certification, you need to have a minimum of at least one year of cumulative paid full-time work experience in one or more of the seven (ISC)² SSCP domains. The domains are as follows:

- Access Controls
- Security Operations and Administration
- Risk Identification, Monitoring, and Analysis
- Incident Response and Recovery
- Cryptography
- Network and Communications Security
- Systems and Application Security



TIP These domains represent the seven major categories of information in the SSCP Common Body of Knowledge (CBK). The CBK is a group of topics updated periodically by subject-matter experts.

When you register for the exam, you identify the number of years' experience that you have in any of the domains. After you take and pass the exam, you're required to submit

a résumé that documents this experience. Additionally, you must find an (ISC)² certified professional (in good standing) to validate your experience and submit an endorsement form on your behalf. You need to complete this endorsement process within nine months of learning that you've passed the exam. (ISC)² completes periodic random audits to ensure the integrity of these documents.

Don't worry if you currently don't know an (ISC)² certified professional. During the course of your studies, you are very likely to meet one or more people who can endorse you. For example, if you join a local security group, you're sure to run into someone with an (ISC)² certification. (ISC)² chapters exist around the world, and you don't have to be certified to join. This page describes the chapters in more detail: <https://www.isc2.org/chapters/default.aspx>. Even if you don't meet an (ISC)² certified professional, you can still get an endorsement through (ISC)². You have to submit a special form to request (ISC)² endorsement (see <https://isc2.org/endorsement.aspx>), and the process takes much longer to get an endorsement than it would take using the normal process.

If you don't have one year of experience, you can still earn the Associate of (ISC)² designation. You still need to complete the other requirements, including passing the exam and subscribing to the Code of Ethics. You'll then have up to two years to obtain the required experience and submit the endorsement form to convert your status from Associate of (ISC)² to become a fully certified SSCP.

Passing the Exam

The exam includes 125 multiple-choice questions, and you'll have three hours to complete it. Most questions are stand-alone multiple-choice questions. However, you may encounter some scenario-type questions that start with a paragraph or two describing the scenario, followed by two or more multiple-choice questions about the scenario. For example, you may be given a scenario that explains an organization's security goals, followed by questions that ask how best to implement those goals.

(ISC)² maintains a large pool of questions and regularly adds new questions to the pool. However, (ISC)² typically tests the new questions on actual exams before using them in the exam score. Up to 25 questions on any exam might be ungraded test questions that (ISC)² has included for research purposes. You can think of these as beta questions that (ISC)² analyzes for the questions' effectiveness. For example, if everyone answers the beta question correctly, (ISC)² may conclude that the question is too easy and needs to be revised. If everyone answers the question incorrectly, (ISC)² will likely conclude that something is wrong with the question itself. Through this type of test question analysis, (ISC)² attempts to identify problematic questions and revise them before including them in the questions that are actually graded. You won't know which questions are graded questions and which questions are beta-type questions. In other words, you have to treat each question as though it's a valid question.

A score of 700 out of a possible 1,000 points is required to pass the exam. However, that doesn't necessarily mean that you'll pass the exam if you answer 70 out of 100 graded questions correctly, because questions aren't weighted the same. Out of the 100 valid questions, some questions may be worth 10 points, more difficult questions may be

worth more than 10 points, and easier questions may be worth less. Just as (ISC)² doesn't tell you which questions are graded, it doesn't advertise the actual value of any question.

(ISC)² derives the examination questions from the SSCP CBK. The SSCP Candidate Information Bulletin (CIB) is an excellent source to see which topics are tested, and it also includes a five-page list of references. The goal of this book is to compile the relevant knowledge from these references into a single source. However, I strongly encourage you to download and review a copy of the current CIB to ensure the document hasn't changed. The CIB includes Key Areas of Knowledge for each of the domains that equate to objectives for the exam. You'll find an Exam Objective Map in the Introduction of this book. It lists the objectives that went live on April 15, 2015, and maps them to the chapter that covers the objectives.



TIP You can request a copy of the SSCP CIB here: <https://www.isc2.org/cib/default.aspx>.

SSCP to CISSP

The (ISC)² Certified Information Systems Security Professional (CISSP) certification is one of the top security certifications. It requires five years of experience in two or more of the eight (ISC)² CISSP domains. It is a highly respected certification and opens many doors of opportunity for those who earn it.

However, not everyone has five years of experience. The good news is that you can request a waiver of one year of experience if you have the (ISC)² SSCP certification. In other words, you only need four years of experience to earn the CISSP if you get the SSCP certification first and receive a waiver.

Many people use the (ISC)² SSCP as a stepping stone to the CISSP. By first achieving the SSCP, you gain an understanding of the (ISC)² certification process. Further, you'll find that the knowledge you gain studying for the SSCP certification helps you to build a solid foundation for the CISSP exam.

Predicting Future Question Types

It's worth noting that (ISC)² has added what it calls “innovative Drag & Drop and Hotspot” questions to the CISSP exam. (ISC)² started using these new types of questions in January 2014, which was two years after the release of the CISSP CIB. (ISC)² didn't change the objectives, but instead changed the way some of the questions are presented. (ISC)² announced the changes a couple of months before implementing them.

(ISC)² has not given any indication that it will implement these types of questions in the SSCP exam. However, it's entirely possible that (ISC)² will start adding them at some point, just as it did with the CISSP exam. If it does so prior to your exam, don't let the questions scare you. If you understand the content, you should be able to answer the questions without too much trouble. The following sections give you an idea of how to handle these types of questions.



TIP I regularly post blog articles on the Get Certified Get Ahead website (<http://blogs.getcertifiedgetahead.com/>). If (ISC)² does announce its intent to add these types of questions, I will post information there to help keep readers informed.

Drag & Drop Questions In a Drag & Drop question, you simply drag an object from one area of the screen to another. As an example, the following question asks you to identify symmetric encryption algorithms from a list:

Question: Which of the following algorithms are examples of symmetric encryption? Drag and drop the correct answers from the list on the left to the box on the right.

Figure 1-1 shows the initial display you might see, along with how the display looks after you have correctly answered the question.

This is really just another way of presenting a multiple-choice question that asks you to select all the correct answers. Here's an example of a multiple-choice question that tests the same knowledge:

Question: Which of the following algorithms are examples of symmetric encryption? (Select all that apply.)

- A. AES
- B. Blowfish
- C. DES
- D. RSA

While you won't see the correct answer and an explanation when you take the exam, you might like this knowledge. For clarity, here is the correct answer and an explanation.

Answer: The correct answers are A, B, and C. Advanced Encryption Standard (AES), Blowfish, and Data Encryption Standard (DES) are all symmetric encryption algorithms.

D is incorrect. RSA (named after its authors Rivest, Shamir, and Adleman) is an asymmetric encryption algorithm.

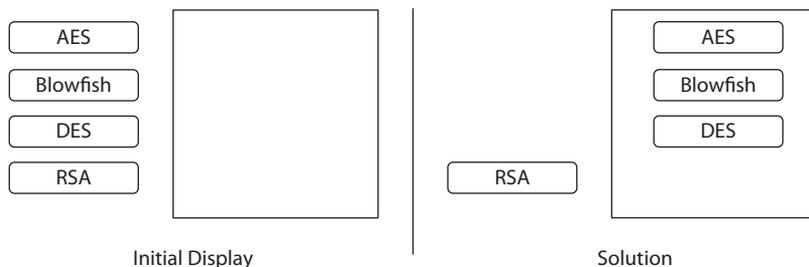
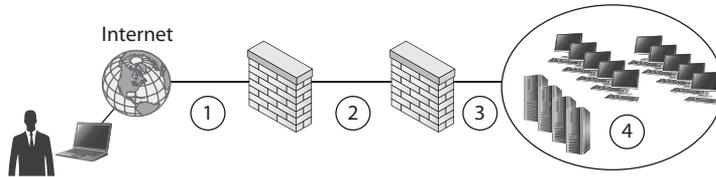


Figure 1-1 Example Drag & Drop question

Figure 1-2
Example Hotspot question



Hotspot Questions In a Hotspot question, you simply click on an area of the screen to select the correct answer. As an example, the following question asks you to identify the best location for a public-facing web server:

Question: Your organization is planning to deploy a web server. It needs to be accessible via the Internet and will access a database on the internal network. Where should they locate the new web server? (Click on the area of the diagram to indicate your choice.)

Figure 1-2 shows the graphics for this question. The numbers (1 through 4) are for the explanation only and probably would not be on an actual Hotspot question.

The Hotspot format is also just another way of asking a question to test your knowledge. Here's an example of a similar question worded as a multiple-choice question:

Question: Your organization is planning to deploy a web server. It needs to be accessible via the Internet and will access a database on the internal network. Where should they locate the new web server?

- A. On the Internet side of the firewalls (the area marked with a 1)
- B. Between the firewalls (the area marked with a 2)
- C. On the private side of both firewalls (the area marked with a 3)
- D. With the database server (the area marked with a 4)

Answer: The correct answer is B. This area is a perimeter network or demilitarized zone (DMZ). It provides a layer of protection against external attacks, and helps prevent a compromised server from accessing internal resources. A is incorrect. Placing the web server directly on the Internet leaves it vulnerable to more attacks. C and D are incorrect. Placing a public-facing server on an internal network increases the risk of attacks for other resources on the internal network.

Maintaining Your SSCP Certification

After you've earned the SSCP certification, you're required to recertify every three years. The primary method of doing this is by acquiring 60 continuing professional education (CPE) credits every three years, with a minimum of 10 CPEs earned each year. Security constantly changes, and earning CPEs is one of the ways security professionals keep abreast of current security trends.

The CPE requirement is a surprise to some people, but many professions use the same concept. As an example, medical doctors are required to complete a minimum number of continuing medical education (CME) credits to maintain their medical licenses.

(ISC)² categorizes CPE credits as Group A credits and Group B credits. Group A credits are for activities directly related to one of the domains in the CIB. Group B credits are optional and are earned for activities that are outside of the domain, but can still enhance a member's general professional skills and competencies.

You typically earn one CPE credit for each hour you spend in a related activity. Group A credits can be earned by attending educational/training courses and seminars, attending conferences, attending vendor presentations, completing some academic courses, and preparing for a presentation, lecture, or training event. Some examples of Group B credits include attending a management course and participating in project planning activities. This is not an exhaustive list, but instead just a few examples.

To maintain your SSCP certification, you must earn the following CPE credits:

- At least 10 Group A CPEs annually
- At least 60 CPE credits during a three-year certification period
- At least 40 Group A CPE credits during a three-year certification period
- As many as 20 Group B CPE credits during a three-year certification period

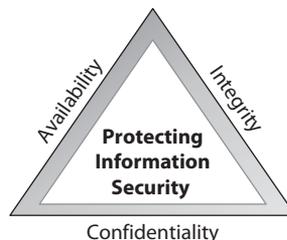


TIP You have a lot of flexibility with Group B credits. Just about anything that enhances your career can be applied as a Group B credit. The only restriction is that you can't use more than 20.

Understanding Basic Security Concepts

Three primary goals of information security are to prevent the loss of confidentiality, the loss of integrity, and the loss of availability for information technology (IT) systems and data. You'll find that most security practices and security controls are designed to help prevent losses in one or more of these areas. The SSCP objectives identify these as the CIA triad, using the initials for confidentiality, integrity, and availability. Figure 1-3 depicts the CIA security triad.

Figure 1-3
CIA security triad
protecting infor-
mation security





NOTE The CIA security triad is sometimes called the AIC security triad. Both are correct because the order of the initials doesn't matter. What you really need to know is what each letter represents (confidentiality, integrity, and availability) and what it means to prevent losses in these areas.

System and information owners are responsible for ensuring that security controls are in place to protect their systems and data. For example, managers that own proprietary data need to ensure that security controls are in place to prevent the unauthorized disclosure of that data. IT security personnel implement and maintain these security controls.



EXAM TIP The three primary goals of an information security program are to prevent the loss of confidentiality, the loss of integrity, or the loss of availability for any IT systems and data. System and information owners are responsible for ensuring that security controls address the confidentiality, integrity, and availability of their systems and data.

Confidentiality

You protect against the loss of confidentiality by ensuring that unauthorized users cannot access data. This starts by identifying and authenticating users, and then implementing access controls to restrict access. For example, you can use permissions to ensure that only authorized users can access the data.



NOTE Chapter 2 covers authentication and access controls in more depth, and Chapter 14 covers encryption algorithms.

Encryption provides another layer of protection for confidentiality. Figure 1-4 shows the overall process of encryption, where data starts as plaintext, is ciphered using an encryption algorithm, and then becomes ciphered text. For example, if you've ever ordered anything over the Internet using a credit card, you've probably used a HyperText Transfer

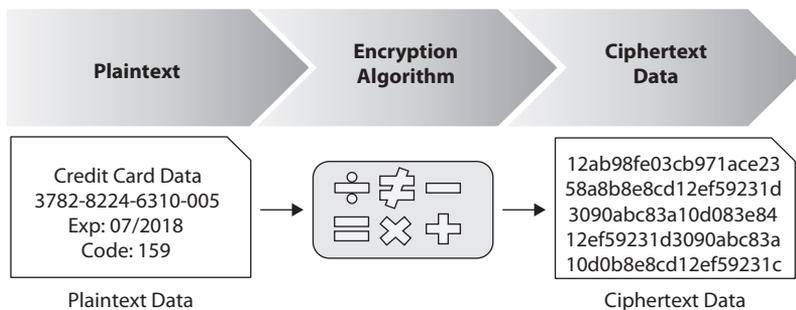


Figure 1-4 Encryption used to provide confidentiality

Protocol Secure (HTTPS) connection. HTTPS encrypted your credit card information to prevent unauthorized individuals from intercepting it and using it without your permission. If criminals had managed to capture this transmission, they wouldn't have been able to read and use your credit card information.

Confidentiality only works when organizations implement secure encryption algorithms and practice sound security practices. With this in mind, it's important to know which algorithms are secure and which algorithms have been compromised and shouldn't be used anymore.

As an example, Wired Equivalent Privacy (WEP), one of the earliest protocols created for wireless transmissions, has significant vulnerabilities, and attackers can crack it using off-the-shelf tools. The Wi-Fi Alliance released Wi-Fi Protected Access (WPA) as an interim replacement for WEP, but researchers later discovered flaws in WPA too. The Institute of Electrical and Electronics Engineers (IEEE) published amendments to the IEEE 802.11i standard (IEEE 802.11i-2004 and IEEE 802.11i-2007), more commonly known as WPA2. At this point, both WEP and WPA are considered compromised and should not be used. Unfortunately, many people are still using these protocols and mistakenly believe that they are protecting the confidentiality of their data.



EXAM TIP Confidentiality controls help prevent the unauthorized disclosure of data. You ensure confidentiality by authenticating users and implementing access controls to ensure that only authorized users can access the data. You can also encrypt data to ensure that even if the data falls into the wrong hands, it's less likely that unauthorized users can read it.

Integrity

Integrity controls prevent any unauthorized or unwanted modification of data or systems. Several different methods are used to protect integrity, including hashing and audit logging. Hashing methods will detect the alteration of data and alert personnel that it has lost its integrity.

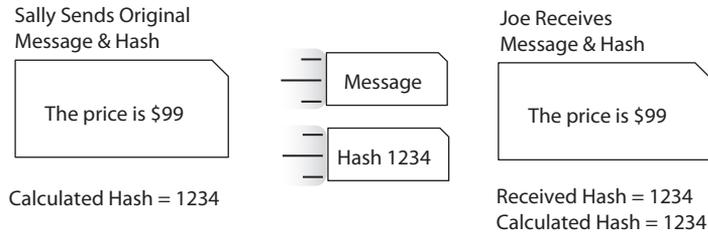


NOTE Chapter 10 covers auditing and logs in more depth, and Chapter 14 covers hashing, including specific hashing algorithms.

A *hash* is simply a number created by performing a mathematical operation against data, such as a file or message. As long as the data stays the same, the hash (the number) will always be the same.

As an example, imagine that Sally needs to inform Joe of the price of a particular company's stock. She creates an e-mail message stating "The price is \$99" and creates a hash of the message. To keep the example simple, the hash for this message is 1234, though an actual hash would be much longer. Sally could re-create the hash on this message 100

Figure 1-5
Hashing used to
check integrity



times and each time it would be 1234. Here's how she can use the hash to validate the integrity of the message after she sends it:

1. Sally sends her message to Joe along with the calculated hash, as shown in Figure 1-5.
2. When Joe's system receives the message, it calculates the hash again.
3. The received hash is the same as the calculated hash, providing assurances that the received message is the same as the sent message.

In contrast, imagine that someone intercepted the message and modified it before resending it to Joe, as shown in Figure 1-6. Joe's system calculates the hash on the received message, "The price is .99," and determines it is 9876. It compares the calculated hash (9876) against the received hash (1234) and discovers that the message is different. In other words, the message has lost integrity.

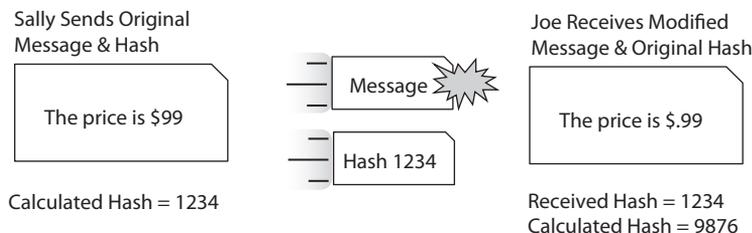


TIP If attackers can change the message, they can also change the hash. Some cryptography systems protect the hash by encrypting it, which prevents attackers from modifying the hash.

Using hashing alone, you can't determine what altered the message. You only know that it was changed. However, this is valuable information. If the message is different, it shouldn't be trusted.

While it's important to be able to verify the integrity of data, integrity also applies to system configuration. As an example, organizations often use change and configuration

Figure 1-6
Hashing identi-
fies altered data



management processes to prevent unexpected system outages resulting from changes to the system. If a technician makes an unauthorized change, it results in a loss of integrity for the system. Worse, many unauthorized changes have caused unexpected system outages.

You can also use audit logging for system integrity. An audit log tracks changes to a resource, including what was changed, who changed it, and when. A set of one or more audit logs creates an audit trail that you can use to verify whether the configuration of a system is the same or has been modified. If someone did make a change that caused an outage, investigators can use audit logs to identify what was changed and who made the change.



EXAM TIP Integrity ensures that data or systems have not been altered. Two common methods used to ensure integrity are hashing and audit logs.

Availability

Preventing the loss of availability ensures that IT systems and data are available when needed. Note that there isn't a timeframe here. Some organizations operate only during the daytime from Monday to Friday, so this is the only time when the systems are needed. Other organizations are operational 24/7, so the systems and data must also be available 24/7.

If users need to access data on a server and they can access it, then the data is available. However, if the data becomes corrupt or the server fails, the result is a loss of availability.

Organizations protect against loss of availability using a variety of different technologies. These include the following:

- **Backups** Regular backups capture a copy of the data. If something happens to the original data, administrators can restore the data from backups. It's important to keep a copy of backup data in an offsite location, so that if a fire or other catastrophe destroys the entire building, the data is still available.
- **Redundant disks** Many Redundant Array of Independent Disks (RAID) systems will continue to operate even if a disk fails. A mirror (RAID 1) is one example of redundant disks. The RAID system stores identical data on two disks, and if one disk drive fails, a copy of the data is still available.
- **Redundant servers** If a service provided by a server is critically important to an organization, the organization can add redundant servers. For example, failover clustering uses multiple servers and ensures that a service will remain available, even if a server fails.



NOTE Chapter 9 covers backups, redundancy, and fault-tolerant techniques in greater depth. Chapter 10 discusses alternate locations used as redundant sites.

- **Redundant connections** Organizations often need to stay connected to the Internet or stay connected between buildings in separate locations via an intranet. When this connectivity is critical to the operation of the organization, two or more connections are used so that even if one fails, the organization still has connectivity.
- **Redundant sites** Many organizations must stay operational even if a catastrophic event destroys their building or makes it uninhabitable. For example, many locations are susceptible to earthquakes, tornadoes, floods, and hurricanes. An organization can plan for these catastrophes by establishing a separate location. Redundant sites are known as *hot sites* (ready at a moment's notice), *cold sites* (an empty building with electricity and running water), and *warm sites* (a cross between a hot site and a cold site).

In addition to using fault-tolerant and redundant technologies, organizations create business continuity plans and disaster recovery plans. These help the organizations maintain the availability of critical systems even after a disaster.



EXAM TIP Availability ensures that authorized users can access any resource when it's needed. Fault-tolerant and redundant technologies ensure that availability is not lost even if a system suffers a failure.

Exploring Fundamentals of Security

In addition to knowing the main goals of security (the CIA security triad), you need to understand some basic terms and concepts for the SSCP exam. The following sections introduce these concepts and some of the common terminology.

Least Privilege

An important security principle is the principle of least privilege. In short, this means that you grant users access to what they need to perform their jobs, and no more. This includes granting permissions to access resources such as files and granting rights to perform actions such as modifying system configurations.

For example, consider a group of project managers and project team members who all need access to a folder named Project Data. The team members need to be able to read the data but not modify it, while the managers need full control over the folder. Table 1-1 shows a simple matrix identifying appropriate permissions. Notice that regular users are not granted any permissions because they don't need access to any of this data to perform their jobs.

If you instead gave the project team members full control permission on the folder, they would still be able to read the files, but they could also modify them or even delete them. Consider what could happen if the team members had full control and one of the team members became a disgruntled employee. That user could deliberately modify or

Table 1-1	User Group	Project Data Folder Permissions
Permission Matrix	Project Managers	Full control
	Project Team Members	Read access (cannot modify)
	Regular Users	None

delete data on the server. Even the most loyal employees can accidentally modify or delete files, but not if they don't have permissions to do so. By assigning only the required permissions, you reduce the risk of anyone modifying or deleting the data.

Similarly, network administrators need elevated privileges to modify network and server configuration settings. Regular users don't need these privileges to perform their jobs, so administrators ensure that regular users don't have them. This reduces the possibility of a regular user accidentally making a change that affects the availability of a system.



EXAM TIP The principle of least privilege ensures that users are granted only the rights and permissions needed to perform their jobs, and no more.

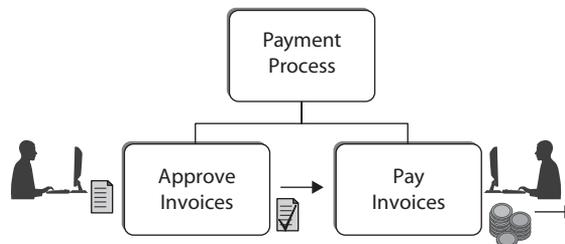
Separation of Duties

Separation of duties is a security principle that ensures that no single person has complete control over a process. When properly implemented, separation of duties significantly reduces the risk of fraud within an organization.

Consider the process of approving and paying invoices. If Joe controlled the entire process, he could create an invoice for his own fictitious company, approve the invoice, and then make a payment to his own bank account. Of course, the loser in this scenario is the company that is employing Joe. Separating the payment process into two steps and assigning different people to handle each step reduces the risk of fraud. Figure 1-7 shows how one person approves the invoice, while another person pays the invoice. Because neither person has full control of the process, neither person can defraud the company without involving the other person.

Implementing separation of duties policies doesn't eliminate the possibility of fraud, because the two employees could choose to collude to defraud the company. Therefore,

Figure 1-7
Separation of duties



many companies also use job rotation and mandatory vacations to reduce the risks of collusion. Chapter 13 covers both of these concepts in more depth.



EXAM TIP Separation of duties helps prevent fraud by ensuring that no single person has complete control over a process.

Privacy

Protecting privacy has become increasingly important within IT security. Two types of data that organizations must take extra steps to protect are personally identifiable information (PII) and protected health information (PHI). PII is information that identifies an individual and includes items such as the person's name, national identification number such as the U.S. Social Security number, and birthdate. PHI is any information about an individual's medical and health history.

Organizations often apply confidentiality principles to protect privacy data such as PII and PHI. This includes using strong access controls to restrict access to the data. It also includes encrypting privacy data as an added layer of protection.



TIP Several laws mandate the protection of an individual's personally identifiable information (PII) and protected health information (PHI). Organizations have a requirement to exercise due care to protect PII and PHI.

Defense in Depth

One of the primary tenets of security is that you're never done. You can't just write a security policy, install antivirus software or enable firewalls, and say, "There. We're safe and secure now." Instead, IT security uses the principle of defense in depth to implement several layers of security.

Consider Figure 1-8. It shows network resources protected through several layers of security. Chapter 9 covers security controls in greater depth, but in short, a security control attempts to reduce risk by either reducing vulnerabilities or minimizing the impact of a threat. One of the primary benefits of a defense in depth strategy is that even if a single control fails, other controls still provide protection.



NOTE The arrows in Figure 1-8 are not meant to imply at which layer any of the controls are implemented. Instead, the message is that IT and security personnel apply multiple methods of security at multiple layers.

For example, you may combine access controls with the principle of least privilege to restrict access to data within your organization. You may also have some research and development data that you want to ensure remains confidential. In addition to access controls and least privilege, you can use cryptography methods to add an extra layer of security for this research and development data. Even if someone is able to bypass the access controls, he or she will not be able to decrypt the data easily.

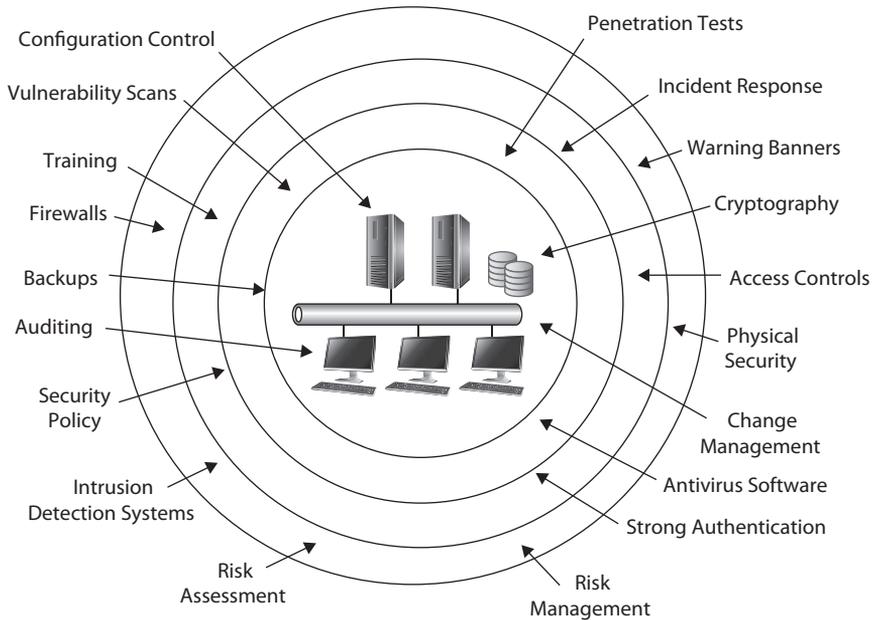


Figure 1-8 Defense in depth includes several layers of security.



EXAM TIP A defense in depth strategy provides a layered approach to security by implementing multiple controls at different layers.

Nonrepudiation

Nonrepudiation ensures that a party cannot believably deny (or repudiate) taking an action. Audit logging and digital signatures are two common methods used to enforce nonrepudiation.

Consider a system that has audit logging enabled for a specific folder. If any user reads, modifies, or deletes data in the folder, the system logs the event in an audit log. The log includes who performed the activity, when they did it, and what they did. If Joe logs on to a computer using his credentials and he deletes a file, the audit log holds a record of his actions. Because the log recorded information from Joe's credentials, you know that Joe did it.



NOTE A remote possibility is that someone else is using Joe's credentials. This possibility increases if the organization uses weak authentication or has poor security practices. However, if Joe logs on with strong authentication (such as with a smart card or biometrics), it's highly unlikely someone is impersonating him.

Digital signatures also provide nonrepudiation. For example, if Sally sends an e-mail to Bob and signs it with a digital signature, Sally can't later deny that she sent the e-mail. Digital signatures use certificates and public/private key encryption. They also provide authentication, giving assurances of who sent the e-mail.

Another example of nonrepudiation is related to commerce and e-commerce transactions. If you use a credit card to purchase a product and you sign the credit card bill, the company can use your signature to prove you are the person who made the purchase. You couldn't later deny it, because your signature verifies that you purchased it. Similarly, e-commerce transactions require you to enter additional information such as the expiration date and the security code on the card. The idea is that only someone with the card in his or her possession knows this additional information.



EXAM TIP Nonrepudiation prevents a party from denying that he or she took an action. The sender of a digitally signed e-mail cannot believably deny sending it. If a system has accountability and an audit trail shows the user took an action, the user cannot believably repudiate it.

AAAs of Security

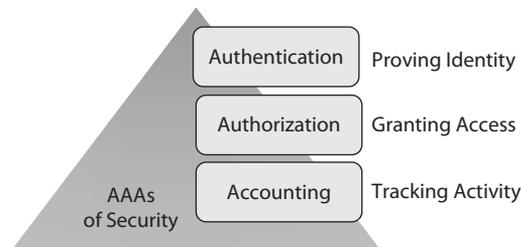
The AAAs of security are authentication, authorization, and accounting. Combined, they help to ensure that only authorized entities have access to resources and that their access is recorded. Figure 1-9 shows the AAAs of security, and the following list explains them. While reading the explanations, imagine that Dawn is a project manager and she needs to modify project files stored on a server.

- **Authentication** A user provides credentials (such as a username and password) that are checked against a database to prove the user's identity. The authentication system verifies the credentials. In the scenario, Dawn's username is her identity, and she provides the correct password to prove her identity. If she enters the correct username and password, the system authenticates her.



EXAM TIP There are three types or factors of authentication, known as (1) something you know (such as a username and password), (2) something you have (such as a smart card), and (3) something you are (using biometrics). Chapter 2 covers authentication in greater depth.

Figure 1-9
AAAs of security



- **Authorization** Administrators assign rights and permissions to resources, which authorize users to access the resources. In the scenario, when Dawn attempts to access files on the server, the system checks to see if Dawn's user account has been granted appropriate permissions and is authorized to access the files. If her account has appropriate permissions, the system grants her access. If her account doesn't have appropriate permissions, the system blocks her access.
- **Accounting** Logging tracks activity of a user through monitoring. A basic accounting mechanism is an audit log, such as the Security log in Windows systems, and audit logs create an audit trail. Imagine that administrators have configured the server storing the project files to log events anytime someone accesses the files. When Dawn opens a file, the system logs an event in the Security log. Similarly, if she saves changes that modify a file, or deletes a file, the system records these events with log entries.

Note that you cannot restrict authorization without proper identification and authentication. If all users have the same account, you can either authorize access to everyone or block access to everyone. Similarly, without authentication, you can't have accounting. If everyone used the same account, event logs could record when users accessed a resource, but they wouldn't be able to record who accessed the resource.



EXAM TIP The AAAs of security are authentication, authorization, and accounting.

Accountability

One of the underlying goals of the AAAs of security is accountability. If a system can identify individual users, track their actions, and monitor their behavior, it provides accountability. Authentication provides identification for users, and accounting tracks their activities in audit logs. If users are not required to authenticate or if audit trails are not created, then a system does not provide accountability.

It's worthwhile noting that you don't have to track every single action of a user to provide accountability. For example, your network may have proprietary data stored in a folder named Research and publicly available information stored in a folder named Public. You may want to track each time any single user accesses any single file within the Research folder. This includes any time someone views, modifies, or deletes a file. However, it's not important to know who viewed any of the data in the Public folder; instead, you only want to know who modified it. In this situation, you would track all activity in the Research folder, but only some of the activity in the Public folder.



EXAM TIP If a system can track activity of an individual on a system, it provides accountability.

Due Diligence

Due diligence refers to the investigative steps that an organization takes prior to taking on something new, such as signing a contract or making a major purchase. In the IT world, an organization has an obligation to exercise due diligence to discover risks associated with a large purchase.

For example, if an organization is planning to purchase a software development company, that organization is obligated to exercise due diligence to determine as much as it can about the company and whether the purchase is a sound decision. Suppose the software company claims it earns about \$100 million in revenue a year, but in actuality only earns \$10,000. If the organization ignores due diligence and simply believes the company's statements without verifying them, purchasing the company will likely result in substantial losses. By exercising due diligence and investigating the claims, the organization will likely discover the discrepancies and avoid the losses.

From an IT security perspective, an organization has a responsibility to exercise due diligence to discover risks. The goal is to identify risks that can result in the loss of availability, integrity, or confidentiality of any data or IT systems. Risks can't be eliminated, but an organization can take steps to mitigate or reduce the risks by exercising due care.



EXAM TIP You cannot eliminate risk. Management decides which risks to mitigate, and the risk that remains is residual risk. Management is responsible for any losses that occur as a result of residual risk.

Due Care

Due care is the practice of implementing security policies and practices to protect resources. It ensures that a certain level of protection is applied to protect against losses from known risks. The goal is to reduce the risk to the resources to a manageable level.

Because risks can't be eliminated, an organization is likely to experience losses. If these losses are due to negligence, then the organization may face legal action against it. However, if the organization took due care to protect the resources but still suffered the loss, it's less likely that a court will find the organization negligent.

For example, imagine that a company holds customer data—including names, addresses, birth dates, and credit card data—in cleartext in a database hosted on a web server. The company uses this information when customers make purchases through a website. A hacker checks out the website, discovers the database, and realizes that he can easily retrieve all the data. He steals the data and sells it to identity thieves, who proceed to steal millions of dollars.

Did the company take due care? Many would say no.

A web server accessible by users on the Internet will be attacked. It's not a matter of *if* it will happen, but *when*. Further, if valuable data is on the server, attackers will find it. Even if it does have some protection, such as encryption, it's still at risk if it's accessible from the Internet.

Taking due care, the organization would implement security controls to protect the data. For example, administrators could store the customer information on a different server within the private network that isn't accessible from the Internet but is accessible

from the web server. Additionally, they could encrypt sensitive data on the server to protect against the loss of confidentiality. They don't need to encrypt the entire database, but instead just encrypt columns holding sensitive data such as credit card data.



EXAM TIP If an organization doesn't take due care to protect private data, it might be sued or forced to pay fines for privacy violations.

Chapter Review

There are several requirements to earn the (ISC)² System Security Certified Practitioner (SSCP) certification. Candidates must have at least one year of experience in one of the seven domains defined by the (ISC)² for the SSCP exam. Prior to registering for the exam, candidates must agree to comply with the (ISC)² Code of Ethics and must answer four questions related to their background.

The SSCP exam is computer based and is administered at Pearson VUE test centers. The exam includes 125 questions, but up to 25 of these are ungraded beta questions. You need to score a minimum of 700 on a scale of 1,000 to pass the exam.

The CIA security triad includes confidentiality, integrity, and availability. System and information owners are responsible for ensuring that security and access controls are in place to prevent losses in these areas.

- Confidentiality protects against the unauthorized disclosure of information. You can ensure confidentiality with access controls and encryption.
- Integrity ensures that unauthorized users have not altered data. Hashing provides assurances that data, such as messages or files, has not been altered. Audit logs provide an audit trail identifying whether data has been modified, who modified it, and when.
- Preventing loss of availability ensures that systems and data are operational and available when needed. Organizations prevent the loss of availability with backups, fault tolerance, and redundant systems.

This chapter also presented several fundamental security terms:

- The principle of least privilege ensures that users are granted only the rights and permissions they need to perform their jobs, and no more.
- Separation of duties helps prevent fraud by ensuring that no single individual has complete control over a process. Instead, tasks are divided between multiple people.
- Organizations are responsible for protecting the confidentiality of privacy data. Many laws mandate the protection of both PII and PHI.
- Defense in depth provides a layered approach to security and protects an organization even if one or more security elements fail.
- Nonrepudiation prevents an individual from denying that he or she took an action. For example, if a user sends a digitally signed message, the user cannot later deny sending it, because a digital signature provides nonrepudiation.

- The AAAs of security are authentication, authorization, and accounting. Authentication identifies users, authorization defines what the users can access, and accounting tracks a user's activities.
- If a system has accountability, it can identify and track the activity of a user. Accountability starts with strong authentication practices.
- Due diligence refers to the investigative steps that an organization takes prior to making a major purchase or taking on an obligation.
- Due care refers to the steps that an organization takes to protect resources. If an organization doesn't take due care to protect resources, such as any individual's private data, it can be susceptible to legal action or fines.

Questions

1. How many years of experience are required to earn the Associate of (ISC)² designation?
 - A. Zero
 - B. One
 - C. Two
 - D. Five
2. What are the three elements of the security triad?
 - A. Authentication, authorization, and accounting
 - B. Confidentiality, integrity, and availability
 - C. Identification, authentication, and authorization
 - D. Confidentiality, integrity, and authorization
3. Who is responsible for ensuring that security controls are in place to protect against the loss of confidentiality, integrity, or availability of their systems and data?
 - A. IT administrators
 - B. System and information owners
 - C. CFO
 - D. Everyone
4. You are sending an e-mail to a business partner that includes proprietary data. You want to ensure that the partner can access the data but that no one else can. What security principle should you apply?
 - A. Authentication
 - B. Availability
 - C. Confidentiality
 - D. Integrity

5. Your organization wants to ensure that attackers are unable to modify data within a database. What security principle is the organization trying to enforce?
 - A. Accountability
 - B. Availability
 - C. Confidentiality
 - D. Integrity
6. An organization wants to ensure that authorized employees are able to access resources during normal business hours. What security principle is the organization trying to enforce?
 - A. Accountability
 - B. Availability
 - C. Integrity
 - D. Confidentiality
7. An organization has created a disaster recovery plan. What security principle is the organization trying to enforce?
 - A. Authentication
 - B. Availability
 - C. Integrity
 - D. Confidentiality
8. Your organization has implemented a least privilege policy. Which of the following choices describes the most likely result of this policy?
 - A. It adds multiple layers of security.
 - B. No single user has full control over any process.
 - C. Users can only access data they need to perform their jobs.
 - D. It prevents users from denying they took an action.
9. Your organization wants to implement policies that will deter fraud by dividing job responsibilities. Which of the following policies should they implement?
 - A. Nonrepudiation
 - B. Least privilege
 - C. Defense in depth
 - D. Separation of duties
10. Which one of the following concepts provides the strongest security?
 - A. Defense in depth
 - B. Nonrepudiation
 - C. Security triad
 - D. AAAs of security

11. Which of the following would a financial institution use to validate an e-commerce transaction?
 - A. Nonrepudiation
 - B. Least privilege
 - C. Authentication
 - D. Signature
12. What are the AAAs of information security?
 - A. Authentication, availability, and authorization
 - B. Accounting, authentication, and availability
 - C. Authentication, authorization, and accounting
 - D. Availability, accountability, and authorization
13. You want to ensure that a system can identify individual users, track their activity, and log their actions. What does this provide?
 - A. Accountability
 - B. Availability
 - C. Authentication
 - D. Authorization
14. Which of the following is required to support accountability?
 - A. Encryption
 - B. Authentication
 - C. Hashing
 - D. Redundant systems
15. Which of the following statements accurately describes due care?
 - A. It is the practice of implementing security policies and procedures to protect resources.
 - B. Due care eliminates risk.
 - C. A company is not responsible for exercising due care over PII.
 - D. Organizations cannot be sued if they fail to exercise due care over resources such as customer data.

Answers

1. **A.** You don't need to meet the experience requirement to earn the Associate of (ISC)² designation, so zero years of experience are required. The SSCP certification requires one year of direct full-time security work experience. If you earn the Associate of (ISC)² designation, you have two years from the date (ISC)² notifies you that you have passed the SSCP exam to obtain the

required experience and apply to become a fully certified SSCP (which includes submitting the required endorsement form). The CISSP certification requires five years of experience.

2. **B.** The CIA security triad includes three fundamental principles of security designed to prevent losses in confidentiality, integrity, and availability. Authentication, authorization, and accounting are the AAAs of security, and identification, authentication, and authorization are required for accountability, but these are not part of the CIA security triad.
3. **B.** System and information owners are responsible for ensuring that these security controls are in place. IT administrators or other IT security personnel might implement and maintain them. While it can be argued that the Chief Executive Officer (CEO) is ultimately responsible for all security, the Chief Financial Officer is responsible for finances, not IT security. Assigning responsibility to everyone results in no one taking responsibility.
4. **C.** Confidentiality helps prevent the unauthorized disclosure of data to unauthorized personnel, and you can enforce it with encryption in this scenario. Authentication allows a user to claim an identity (such as with a username) and prove the identity (such as with a password). Availability ensures that data is available when needed. Integrity ensures that the data hasn't been modified.
5. **D.** Integrity ensures that data is not modified, and this includes data within a database. Accountability ensures that systems identify users, track their actions, and monitor their behavior. Availability ensures that IT systems and data are available when needed. Confidentiality protects against the unauthorized disclosure of data.
6. **B.** Availability ensures that IT systems and data are available when needed, such as during normal business hours. Accountability ensures that users are accurately identified and authenticated, and their actions are tracked with logs. Integrity ensures that data is not modified. Confidentiality protects the unauthorized disclosure of data to unauthorized users.
7. **B.** Availability ensures that IT systems and data are available when needed. Disaster recovery plans help an organization ensure availability of critical systems after a disaster. Users prove their identity with authentication. Integrity provides assurances that data and systems have not been modified. Confidentiality protects against the unauthorized disclosure of data.
8. **C.** The principle of least privilege ensures that users have access to the data they need to perform their jobs, but no more. Defense in depth ensures an organization has multiple layers of security. Separation of duties ensures that no single user has full control over any process. Nonrepudiation prevents users from denying they took an action.

9. **D.** Separation of duties helps prevent fraud by dividing job responsibilities and ensuring that no single person has complete control over an entire process. Nonrepudiation ensures that parties are not able to deny taking an action. The principle of least privilege ensures that users have only the rights and permissions they need to perform their jobs, but no more. Defense in depth provides a layered approach to security.
10. **A.** Defense in depth provides a layered approach to security by implementing several different security practices simultaneously and is the best choice of the available answers to provide the strongest security. The security triad (confidentiality, integrity, and availability) identifies the main goals of security. Nonrepudiation prevents an individual from denying that he or she took an action. The AAAs of security are authentication, authorization, and accounting.
11. **A.** Digital signatures used by some online institutions to validate transactions and provide nonrepudiation. Least privilege ensures that users have only the rights and permissions they need to perform their jobs, and no more. Authentication verifies a user's identity. A written signature is not used in e-commerce.
12. **C.** The AAAs of information security are authentication, authorization, and accounting. Availability is part of the CIA security triad (confidentiality, integrity, and availability), but it is not part of the AAAs of information security.
13. **A.** If a system can identify individual users, track their activity, and log their actions, it provides accountability. Availability ensures the system is operational when needed. Authentication identifies the individual using credentials. Authorization identifies resources that a user can access.
14. **B.** Users prove their identity with authentication, and strong authentication mechanisms are required to support accountability. Encryption helps provide confidentiality. Hashing helps provide integrity. Redundant systems help provide availability.
15. **A.** Due care is the practice of implementing security policies and procedures to protect resources. You cannot eliminate risk. A company is responsible for exercising due care over PII and can be sued if it fails to exercise due care.

This page intentionally left blank

Access Controls

In this chapter, you will learn about

- Identification, authentication, and authorization
- Three factors of authentication
- Single sign-on
- One-time passwords
- Access controls with subjects and objects
- Access control models (DAC, non-DAC, MAC, RBAC, and ABAC)
- Bell-LaPadula, Biba, Clark-Wilson, and Chinese Wall architectures
- Identity management
- Physical security controls

Comparing Identification, Authentication, and Authorization

A key part of any access control system is the identification and authentication of individuals. If you can't identify individuals, everyone is anonymous. If everyone is anonymous, there is no way you can control access to different resources. Either everyone has access or no one does.

However, if you are able to differentiate between different users, you can grant access to some users while denying access to other users. This process starts with identification. In many authentication systems, the identity of a user is simply the user's name.

Figure 2-1 shows the overall process of identification, authentication, and authorization. The user professes an identity with a username and validates the identity by providing authentication information. An authentication system verifies the credentials. Then access controls (such as permissions) authorize the user to access resources such as files stored on a server. If these three steps don't come together, the user is not able to access resources protected with access controls.

Just because someone knows a username doesn't mean they are authenticated. They must also prove their identity by providing authentication data, such as a password or fingerprint. Similarly, just because a user can log on with a username and password doesn't mean they can access everything on a system or a network. Administrators authorize access by configuring rights and permissions to users based on their proven identities.

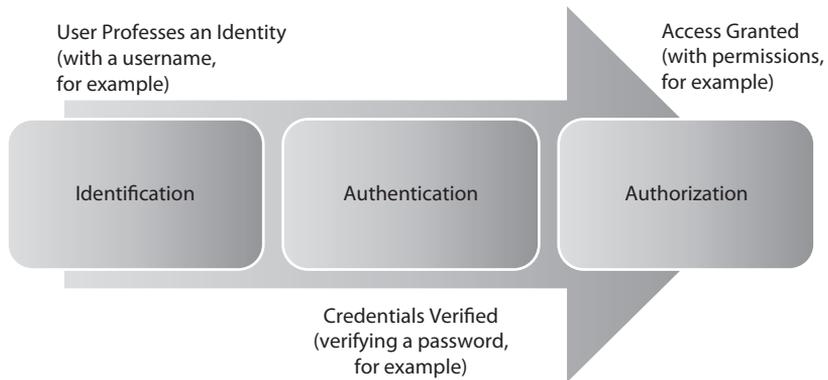


Figure 2-1 Identification, authentication, and authorization



EXAM TIP Identification and authentication are two primary controls used by access control systems. The user professes an identity and then proves his or her identity with authentication. Administrators grant a user rights and permissions based on the user's proven identity.

Exploring Authentication

Authentication is the process of users proving they are who they claim to be. A common method is by having a user provide a logon name to profess an identity and then using a password for authentication.

For example, Figure 2-2 shows a challenge screen where a user must provide account credentials. The username provides the identification, and the password provides the authentication.

While you've probably used a password to authenticate to a system, you might not have used all possible authentication mechanisms. The following sections cover the different authentication methods.

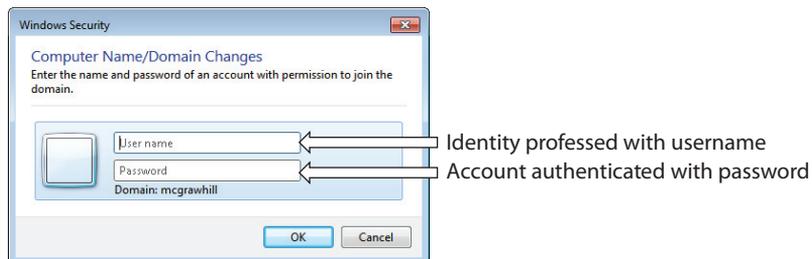


Figure 2-2 Username and password for identification and authentication

Three Factors of Authentication

As shown in Figure 2-3, most authentication falls into one of three primary types or factors:

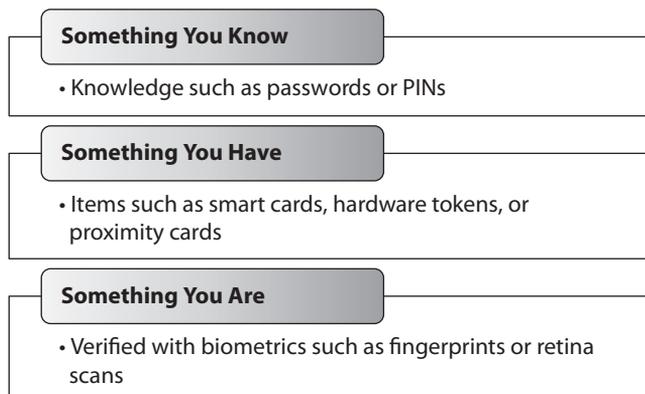
- **Something you know (type 1)** This includes knowledge, such as passwords, personal identification numbers (PINs), your mother’s maiden name, or even personal information such as the name of your first pet.
- **Something you have (type 2)** This includes items such as smart cards, hardware and software tokens, and proximity cards.
- **Something you are (type 3)** This includes the use of biometrics, such as fingerprints and retina scans.

Something You Know

The simplest type of authentication is *something you know* and includes passwords, PINs, and other information known by an individual. Passwords have the following classifications:

- **Static password** A static password stays the same over a period of time. Each time a user logs on to a session, they use the same password.
- **One-time or dynamic password** One-time passwords are used only once per session. Many hardware tokens generate new passwords every 60 seconds that users can enter as a one-time password. The “Something You Have” section, later in this chapter, explains hardware tokens in more depth.
- **Cognitive password** Information that a user would know, such as the name of a first pet, a first employer, and a favorite color, is combined to create a cognitive password. Users enter this information during a registration process. Later, the authentication system challenges the user with these questions in specific situations. For example, many authentication systems register information about the user’s computer, and if they detect the user is accessing the

Figure 2-3
Three factors of authentication



account from a different system, they challenge the user with cognitive password questions. As long as the user can give the same information when accessing the account at a different computer, the system authenticates the user.

- **Passphrase** A passphrase is a long string of characters that has meaning to the user. Users typically code these to make them more complex. For example, you can create a passphrase of *IWi11P@\$\$The\$\$CPEExam* from the statement *I will pass the SSCP exam*. It's a long password of 20 characters, but relatively easy to remember. The letter *l* is replaced by the number 1 in each instance, the *a* is replaced by @, each *s* or *S* is replaced by \$, and the first letter in every word is capitalized. Even as complex as it is, most people could easily commit this 20-character password to memory.

Passwords are the least secure method of authentication; however, you can increase password security by following some basic guidelines:

- **Use strong passwords** A strong password includes a combination of different character types, such as uppercase letters, lowercase letters, numbers, and symbols. A strong password also needs to be sufficiently long. While security professionals previously recommended using passwords that are at least 8 characters long, many professionals now recommend using passwords at least 12 characters long. Strong passwords do not include usernames, logon names, children names, pet names, or any words found in a dictionary (of any language).



NOTE Some systems create random passwords for users. However, these are often much more difficult for users to remember, causing them to write them down.

- **Don't write passwords down** If a password is only in a person's head, it can't be read. On the other hand, if a user writes it down on a yellow sticky note and sticks it beneath a keyboard, an attacker can find it.
- **Change passwords often** Passwords should be changed at least every 90 days, although many organizations require users to change their passwords more often, such as every 30, 45, or 60 days.
- **Don't use the same password on multiple systems** Unfortunately, attackers have been successful at hacking into systems and downloading huge databases that include user credentials. Attackers then try these credentials on other sites. If users have the same credentials on multiple sites, it becomes easier for attackers to compromise their accounts.
- **Never give your password out** Many social engineering attacks can easily be avoided if users understand and follow this simple rule. Because social engineers frequently try to get passwords simply by requesting them, security professionals repeat this rule to end users often.



NOTE Chapter 5 covers social engineering and many types of social engineering attacks.

- **Audit passwords** Many organizations use technology such as built-in operating system tools or third-party applications to verify that passwords are strong and that users change them regularly. A password audit ensures that passwords are strong (never blank), are of a minimum length, and have been changed within a given time period (such as 90 days). This is much more effective than a security professional trying to verify that all user passwords are strong and changed regularly.
- **Use a credential management system** Credential management systems provide a storage space for users to keep their credentials. For example, Windows systems include the Credential Manager tool, which allows users to store different username/password pairs and certificates used for authentication. Similarly, there are many password wallet or password vault apps that allow users to store multiple passwords within a single application. For example, KeePass is a free open source utility that allows users to store usernames and passwords within a secure encrypted database. Users open the database by entering a single master password or key file.



EXAM TIP Some password audit tools only enforce the use of strong passwords when users create the password. Other tools can also check passwords periodically to ensure that password policies are followed, such as ensuring that they have been changed within a given time period.

Many organizations implement written password policies providing rules for users on passwords. The policy identifies the minimum requirements. Administrators enforce the password policy using a technical control. For example, Windows Group Policy includes Password Policy settings that administrators can configure to apply to all users in a domain.

Figure 2-4 shows the Windows Local Group Policy Editor with the Password Policy selected. The following list describes the Password Policy settings:

- **Enforce password history** This setting remembers previous passwords that a user has created and prevents the user from reusing the same passwords. As configured, it prevents a user from reusing any of the last 12 passwords they've created.
- **Maximum password age** This setting defines when users must change their password. As configured, it requires users to reset their password every 30 days.
- **Minimum password age** This setting defines how long users must wait before changing their password again. As configured, it prevents users from changing

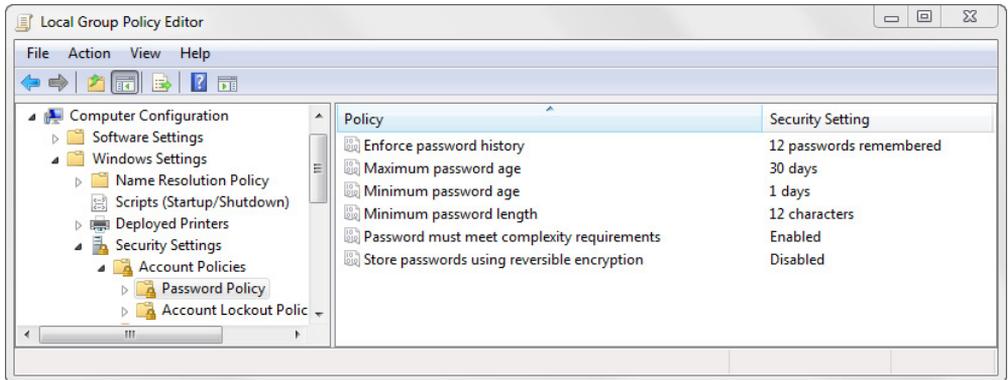


Figure 2-4 Password Policy within Windows Local Group Policy

their password until 1 day has passed. This setting works with the password history setting. If set to 0, users can change their password 12 times in a couple of minutes and get back to their original password.

- **Minimum password length** This setting enforces the character length of the password. As configured, it requires users to create passwords at least 12 characters long.
- **Password must meet complexity requirements** When enabled, this setting requires users to create complex passwords that include at least three of the four character types—uppercase letters, lowercase letters, numbers, and special characters.
- **Store passwords using reversible encryption** Some legacy systems require storing passwords with reversible encryption. However, this is not recommended because it makes it easier for attackers to discover the password. This setting should be disabled unless otherwise required.



EXAM TIP A written password policy often identifies requirements for strong passwords, the minimum amount of time before a password should be changed, and how passwords are audited. Technical controls enforce the policy. Many technical password policies can remember previous passwords and prevent users from reusing the same passwords again.

A popular method to create strong passwords is to use a password manager such as KeePass. It will create and store complex passwords within a database. As long as you use a strong master password to open the application, it keeps your passwords relatively secure. This is much better than storing all your passwords in a file named passwords.

How Strong Is Your Password?

Passwords take significantly longer to crack as you add more characters and as you add complexity. One site (<http://howsecureismypassword.net/>) calculates how secure your password is by letting you know how long it would take a system to crack it. When you enter a password, the site tells you about how long it takes a desktop PC to crack your password. Although the times may not be exact, you can see the big difference in times when you add more characters to the password.

- Five characters takes less than a second to crack.
- Six characters with a mixture of three character types takes about 14 seconds.
- Six characters with a mixture of four character types takes about 52 seconds.
- Seven characters with a mixture of four character types takes about an hour.
- Eight characters with a mixture of four character types takes about three days.
- Nine characters with a mixture of four character types takes about 275 days.
- Ten characters with a mixture of four character types takes about 58 years.
- Eleven characters with a mixture of four character types takes about 4,000 years.
- Twelve characters with a mixture of four character types takes about 344,000 years.

Note that these timeframes are a moving target. Computer processing power doubles every 18 to 24 months, allowing desktop PCs to crack these passwords quicker and quicker.

Something You Have

Organizations that want to increase the security of authentication often use the next factor: *something you have*. Smart cards, hardware and software tokens, and proximity cards are some examples.

Smart Cards A smart card is a small card about the size of a credit card that has an embedded certificate used to identify the user. The user inserts the card into a smart card reader that reads the data on the card to identify and authenticate the user. Because it's easy for a user to lose a smart card, most smart card authentication systems require users to enter a PIN or a password in addition to using the smart card.



EXAM TIP Combining the something you have factor with the something you know factor provides two-factor authentication. For example, when a user logs on with both a smart card and a PIN to authenticate, the user is authenticating with two-factor authentication. Two-factor authentication is more secure than just a single factor of authentication.



Figure 2-5 Token used for token-based authentication

Hardware Tokens and One-time Passwords A hardware token is a small device that displays a number. The number changes periodically, such as every 60 seconds, and the user enters the number for authentication. This number is a *one-time password* because the user can only enter it once. You'll often see this called a *dynamic password* to contrast it with a static password that stays the same. A one-time password reduces the risk of an attacker capturing the password and reusing it.

Authentication servers know the number displayed on the token, as shown in Figure 2-5. A user typically authenticates via a website by entering the username, password, and the number displayed in the token, which is sent to the authentication server. If the user enters the correct number, it proves that the user has the token. As long as the user entered the correct username and password, this provides authentication. As with smart cards, hardware tokens are typically used with another factor of authentication. In this example, the user entered a password and the number displayed on the token.

Hardware tokens can use either a synchronous dynamic password or an asynchronous dynamic password:

- **Synchronous dynamic password** A token using a *synchronous dynamic password* changes the password at specific times, such as every 60 seconds. This requires the token and the server to be synchronized with the same time.
- **Asynchronous dynamic password** A token using an *asynchronous dynamic password* doesn't require exact clock synchronization. Instead, it uses another method to create one-time dynamic passwords. As an example, the user can enter a PIN into the device, which prompts the token to generate a one-time password.

Software Tokens and One-time Passwords Software tokens are being used more and more today. These use open source protocols such as the open-standard Time-based One-Time Password (TOTP) protocol or the HMAC-based One-Time Password (HOTP) protocol to create one-time passwords. Google Authenticator and Symantec's VIP Access are two examples of software tokens.

The TOTP protocol creates synchronous dynamic passwords based on time. The HOTP protocol creates asynchronous dynamic passwords. Instead of using times, HOTP uses an incrementing counter combined with a secret key, known by both the server and the token. HOTP creates a hash using the hash message authentication code (HMAC)

and then reduces this hash to a six- to eight-digit HOTP value. Because HOTP uses an incrementing counter, these one-time passwords remain valid until used.

One-time Password In Everything (OPIE) is another type of one-time password used in networks. OPIE is based on S/KEY, a one-time password system used on some UNIX systems. The goal of both OPIE and S/KEY is to generate an asynchronous password known to both a client and a server and used in a single session. Both S/KEY and OPIE combine the user's real password with other data, and use a hashing algorithm such as Message Digest 5 (MD5) to create a one-time password.



EXAM TIP Tokens synchronized with a server based on time create synchronous dynamic one-time passwords. Asynchronous dynamic passwords are not based on time. HOTP, OPIE, and S/KEY are some examples of protocols that create asynchronous dynamic passwords.

Proximity Cards A proximity card includes data electronically embedded within the card. It doesn't have power, but does have an inductor that can be electrically excited when it passes through a magnetic field. This provides enough power for the proximity card to transmit data to a nearby receiver. For example, some credit cards have this technology embedded within them. When you pass the card over the proximity card reader, it activates the card's electronics and transmits your data to the card reader.

Many organizations issue identification badges to employees, which they show to security guards for identification, and some identification badges include proximity card capabilities. For example, an employee can wear the badge for identification and also use it to open doors by passing the badge in front of a proximity card reader.

Something You Are

Biometrics are used to provide authentication as *something you are* by identifying unique characteristics of a person, such as a fingerprint. Users register the biometric value and associate it with an identity such as a username. Later, users authenticate by entering the username and presenting the biometric value. For example, a user can place their finger in a fingerprint scanner. The following list shows some of the biometric methods used for authentication:

- **Fingerprints and thumbprints** Fingers and thumbs have unique patterns that law enforcement personnel have used for decades to identify individuals. Many IT systems now use fingerprints and thumbprints for authentication. As an example, many laptops and small devices include a fingerprint scanner. However, some fingerprint scanners can be tricked with a picture of a fingerprint or even a duplicate fingerprint on a gummy bear.
- **Palms** A palm scanner can measure the vein pattern in a person's palm using an infrared scanner. Users don't need to touch the scanner, but instead just hover their hand over the scanner.

- **Retina** A person's eyes have a pattern of blood vessels at the back of the eye, and a retina scan uses an infrared light to measure that pattern. While retina scans are very accurate, people object to them because they can reveal medical conditions. Also, retina scans typically require physical contact with the scanner.
- **Iris** The area surrounding the eye's pupil is the iris, which is almost as unique as the retina. Iris scans are more acceptable to users because they don't require physical contact and cameras can take pictures of the iris from a distance. However, lighting can affect the accuracy of an iris scan, and some iris scanners can be tricked with a high-quality picture.

Behavioral biometrics identify behavioral traits of an individual. For example, key-stroke dynamics measures the pattern and rhythm as a user types on a keyboard. Similarly, handwriting analysis can sometimes verify that a specific person wrote a note. However, these methods aren't as reliable for authentication as biometric methods that focus on a person's physical characteristics.

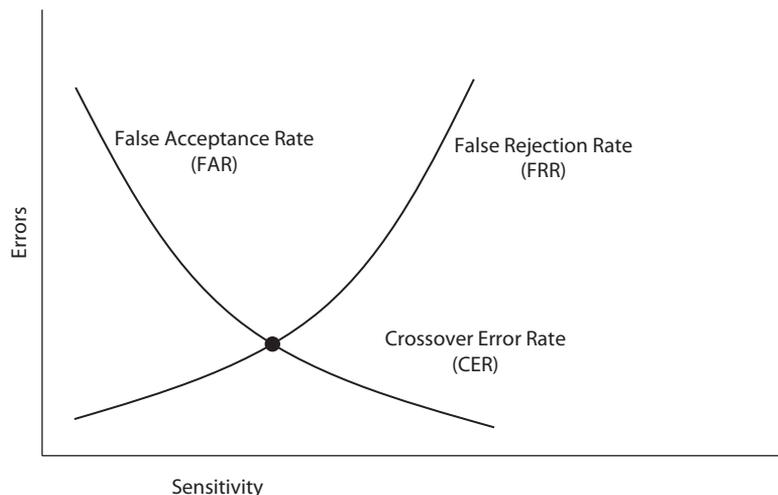


NOTE While this section focuses on using biometrics for authentication, biometrics can also be used for identification. For example, law enforcement has used fingerprint technologies for more than 100 years for accurate identification of criminals because fingerprints are unique for each person.

One of the challenges with biometrics is the potential for errors. When considering a biometric authentication system, it's important to understand the different error types and the accuracy of the system. Refer to Figure 2-6 while reading the following error rate definitions:

- **False Rejection Rate (FRR)** Also called a type 1 error, FRR refers to the percentage of times a biometric system falsely rejects a known user. Instead, the system indicates that the user is unknown.

Figure 2-6
Biometric
error rates



Biometric Errors

Ideally, a biometric system won't have any errors at all, but the cost of such a system is prohibitive for most organizations. Instead, an organization often weighs what it considers more important and adjusts the system to accept more of one type of an error than another error.

For example, consider a Top Secret military facility that chooses to use biometrics for authentication. The most important result here is that the system does not authenticate unauthorized personnel. In other words, the FAR must be as close to zero as possible. However, it is acceptable to have a higher FRR where the system occasionally rejects known users.

- **False Acceptance Rate (FAR)** Also called a type 2 error, FAR refers to the percentage of times a biometric system falsely identifies an unknown user. Instead, the system indicates the user is a known user.
- **Crossover Error Rate (CER)** Also called the Equal Error Rate (EER), CER identifies the point where the FAR and FRR of a biometric system are equal or cross over each other on the chart. A lower CER indicates a better-performing biometric system.



EXAM TIP The three primary performance measurements used in biometrics are the False Acceptance Rate (FAR), the False Rejection Rate (FRR), and the Crossover Error Rate (CER). Systems with low CERs are more accurate than systems with high CERs.

Multifactor Authentication

Multifactor authentication uses two or more factors of authentication. This can be any combination of two of the three factors (something you know, something you have, and something you are) or a combination of all three factors.

Some examples of multifactor authentication include the following:

- A smart card (something you have) and a PIN (something you know)
- A fingerprint (something you are) and a password (something you know)
- A hardware token (something you have) with a username and password (something you know)

A key point is that the different authentication methods must use different factors. For example, using a password and a PIN is not multifactor authentication because it uses two authentication methods of the same type, something you know.



EXAM TIP Multifactor authentication requires users to use more than one factor of authentication. When two factors are required, it's also called two-factor authentication. Multifactor authentication is more secure than any single authentication type used by itself.

Reviewing Identification

A username is a primary method of identification, but it's important to realize that there are more methods. Further, many of these methods are used for both identification and authentication.

Biometric systems are effective at accurately identifying individuals using a one-to-many search. As an example, the TV show *Las Vegas* showed some of the security technologies used by casinos, including face-recognition technologies. Cameras capture the facial characteristics of a person and then search a database of many faces to find a match. Similarly, many passport-free border crossings use iris scans to identify individuals.

In contrast, biometric systems used for authentication use a one-to-one match. For example, if a user claims an identity with a username and then authenticates with a fingerprint, the authentication system compares the user's fingerprint against the fingerprint associated with the user.

Smart cards can also be used for identification. A user inserts the smart card into a reader for identification and then enters a PIN for authentication. For example, the U.S. government uses Common Access Cards (CACs) and Personal Identity Verification (PIV) cards, which are specialized types of smart cards. They include a photo, which users can also present for identification.

Single Sign-on Authentication

With single sign-on (SSO), a user authenticates once and then the system uses the same credentials for the entire session. This includes when the user accesses any resource in the organization, or with a trusted partner. This increases security because users have to remember only a single set of credentials and are less likely to write them down.

Consider the alternative in an organization that has seven servers that any user can access but where each server requires a different username and password. Users have to memorize seven sets of credentials. Even if the usernames are the same on each server, it is almost impossible to ensure that the passwords are the same on all systems unless the passwords are never changed.

However, if the organization implements an SSO system, users would need to remember only a single username and password. This provides convenience for the users. Additionally, with fewer centralized user accounts, it decreases the administrative workload. For example, it is easier for administrators to implement password policies with a centralized SSO system. Of course, the major drawback of an SSO system is that if an attacker obtains a user's credentials, the attacker can access multiple systems. The following sections discuss some commonly used SSO systems.

Kerberos

Windows domains and Linux/UNIX realms use Kerberos as the underlying protocol for SSO. A user can log on to the domain at the beginning of the workday, and the system uses the same credentials to access any resources in the domain.

Kerberos is a vendor-neutral authentication protocol originally developed at the Massachusetts Institute of Technology (MIT) for UNIX realms. It uses a complex process of issuing tickets to authenticated accounts and then uses the tickets to access resources. Microsoft has adapted Kerberos for its Windows domains since Windows 2000.

The Kerberos server issues different types of tickets used for authentication, and is often referred to as a ticket-granting server. Kerberos uses symmetric keys and symmetric encryption (also called secret key cryptography) to protect the confidentiality of the Kerberos transmissions, and you'll also see the Kerberos server referred to as a Key Distribution Center (KDC).

The following steps and Figure 2-7 provide an overview of how Kerberos works:

1. The client (the user's computer) sends an authentication request to the Kerberos server. This includes the user's login ID or username, but not the user's password.
2. After verifying the user's login ID with a database (such as Microsoft Active Directory), the Kerberos server creates a symmetric key that will be known only to the client and the Kerberos server. It encrypts this symmetric key using a hash of the user's password. It also encrypts a time-stamped ticket-granting ticket

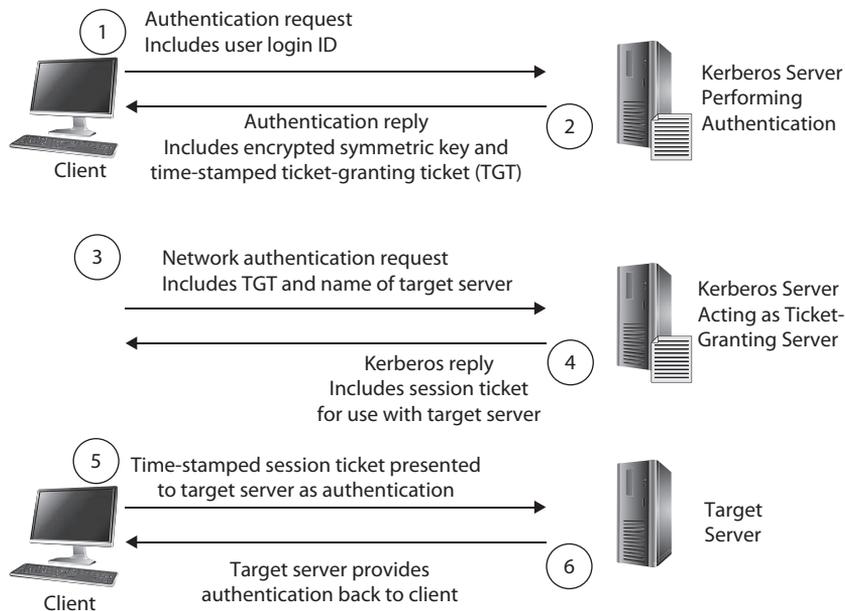


Figure 2-7 Kerberos ticket-granting process

(TGT) with a separate key, and sends both the encrypted symmetric key and encrypted TGT to the client.

The client decrypts the symmetric key with a hash of its password. It can then use this symmetric key to encrypt traffic between the client and server. Note that this verifies the password too. If the client doesn't know the correct password, it isn't able to decrypt the symmetric key.

3. When the client wants to access resources on a target server (such as a file server on the network), it sends the encrypted TGT to the ticket-granting server with the name of the target server.
4. The ticket-granting server replies with an encrypted ticket (also called a session ticket) that the client uses to communicate with the target server. The ticket-granting server also provides an encrypted symmetric key that systems use to encrypt information between the client and the target server.
5. The client presents the encrypted time-stamped session ticket to the target server for authentication.
6. The target server decrypts the session ticket and verifies it. It then replies with a message, which includes the session ticket's time stamp. This authenticates the server to the client (providing mutual authentication), because other servers would not be able to decrypt the session ticket.

Kerberos also provides a level of confidentiality and integrity for information. It helps prevent the unauthorized access to data by ensuring that only authenticated users can access data on target servers. This also contributes to integrity by preventing unauthorized modifications of data on target servers.



EXAM TIP Kerberos provides authentication on a network and contributes to the confidentiality and integrity of information. It uses symmetric encryption (also called secret key cryptography) to encrypt tickets in a secure format between systems. Kerberos requires a database of accounts and time synchronization for all systems.

Federated Access

Many SSO systems use federated access technologies. Federated access allows users in different networks to log on only once, even if they are accessing multiple systems. The systems can be different operating systems, owned and managed by different organizations.

For example, a company can have an intranet website used by employees that also includes links to outside websites. The outside websites can be for employee financial services, such as 401K information, or insurance company websites used to provide information on an employee's health plan.

Users log in to their work account using their regular account. When they access the intranet website, it uses typical SSO technologies such as Kerberos to grant them access. Then if they click on any links to the external sites, the federated access SSO system verifies the user's identity and provides access without additional user interaction.

Each of the different websites has its own independence or autonomy. In other words, the owners of the financial website are not dependent on information on the intranet website or the health insurance website. However, the different websites do provide information on the identity of users and provide assurances to the other sites that the user has authenticated.

Many federated SSO systems share information via a federated database. A federated database does not include passwords of users, but it does include identity information needed by each of the sites. It provides the users with a federated identity, which each of the sites can use.



EXAM TIP Federated access provides SSO to different operating systems or networks. A federated database provides central SSO authentication.

SAML

Security Assertion Markup Language (SAML) is an Extensible Markup Language (XML)-based data format used for SSO on the Internet. As an example, consider two websites hosted by two separate organizations. Normally, a user would have to log on to each site separately. However, the organizations can use SAML as a federated identity management system. Users authenticate once with the first website, and they are not required to authenticate again when they access the second website.

Many online banking sites use SAML for SSO. For example, the banking site might have one service for accessing checking and savings accounts, another service for online bill paying, and another service for handling mortgages. With SSO, the user is able to log on to the primary banking site one time, and then access all the services without logging on again.

SAML includes three roles:

- **Principal** The principal is typically the user that logs on to the system. If necessary, the user might request a principal identity from the identity provider.
- **Identity provider** The identity provider creates, maintains, and manages the identity information for principals.
- **Service provider** A service provider is the entity that provides services to principals. For example, a banking institution that hosts different banking services is the service provider.

The primary purpose of SSO is for identification and authentication of users. Users claim an identity and prove that identity with credentials, such as a password. SSO does not provide authorization. For example, if a user logs on to the banking site, SSO doesn't control what services the user can access. However, SAML includes the ability to transfer authorization data between the bank's systems. In other words, it's possible to use SAML for SSO authentication and for authorization.



EXAM TIP SAML provides SSO for web-based application servers on the Internet. It is an XML-based standard used to exchange authentication and authorization information between different parties.

Other SSO Technologies

Two other SSO technologies you might run across are SESAME and KryptoKnight:

- **Secure European System for Applications in a Multivendor Environment (SESAME)** SESAME was created as an alternative to Kerberos in European countries. However, with improvements to Kerberos, SESAME is rarely, if ever, used today.
- **KryptoKnight** IBM created KryptoKnight as an alternative to Kerberos. It does not have as much network overhead as Kerberos. However, like SESAME, KryptoKnight is rarely used today.



EXAM TIP Kerberos, SAML, SESAME, and KryptoKnight each provide SSO capabilities. With SSO, users only have to log on once and then use the same credentials to access multiple resources.

Centralized vs. Decentralized Authentication

Authentication can be decentralized or centralized. Consider Figure 2-8. The group of users on the left is in a decentralized environment. Every computer has a separate database that stores credentials. If Sally needed to log on to all four computers in this network, she would need to have four separate sets of credentials—one for each system.

On the other hand, the computers on the right are using centralized authentication. Credentials for the users are stored on a central server. Any user is able to log on to the

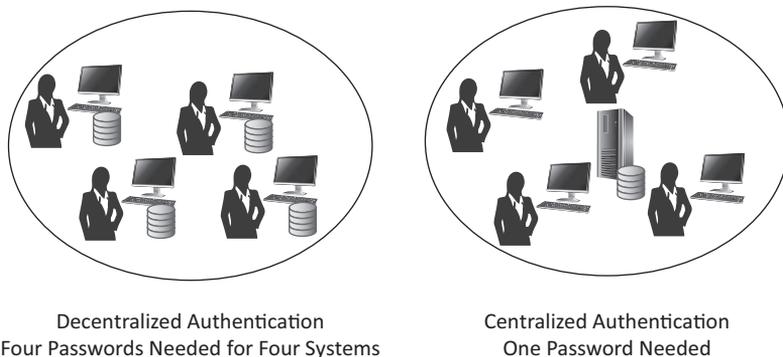


Figure 2-8 Centralized vs. decentralized authentication

network once and then access any computer in the network (as long as the user has permissions). For example, if a computer is part of a Microsoft domain, the central server will be a domain controller and hold accounts for all users in the domain.

Centralized authentication is used for single sign-on networks. It's not possible to use SSO without some type of centralized authentication.

Offline Authentication

Offline authentication allows users who have logged on to the system at one time to still log on even when they are disconnected from a network.

For example, consider a consultant named Steve with an account in a Windows domain named mcgrawhill. He uses his laptop computer to log on to the domain while he's at work. When his company deploys him to a customer site to provide consulting services, he's no longer connected to the mcgrawhill domain, so there's no way that a Kerberos server in the domain can authenticate him.

However, he is still able to log on while he's offline. In a Windows environment, the system uses cached credentials. In other words, when Steve logged on to his laptop, encrypted data about his credentials were stored (cached) on his system. If he's disconnected when he logs on, these cached credentials are compared to the credentials he's entering, and if they match, he is authenticated.



TIP A user will not be able to access network resources while using cached credentials. The user can only access resources on the local system.

Linux and UNIX users have a similar capability when they are interacting with a Microsoft domain. As long as offline authentication is selected when the computer joins the domain, they will be able to use domain credentials even when they are disconnected.

Device Authentication

Device authentication methods focus on hardware instead of the user. The goal is to prevent unauthorized devices from accessing the network. This has become increasingly important to organizations as users try to connect their mobile devices (such as smartphones and tablets) to the network. Some organizations embrace bring your own device (BYOD) policies, while other organizations attempt to block this activity due to the risks.

Windows domains require users to log on from known computers that have joined the domain. This requires the computer and the user to authenticate before the user is able to access domain resources. However, there are times when you might want to allow network access without joining the device to the domain.

A traditional method used for device authentication allows or blocks devices with Media Access Control (MAC) address filtering. However, this is very tedious and users can easily bypass MAC address filtering if they know how to change their MAC address.

Another method is device fingerprinting that identifies devices based on several characteristics. Many of the following characteristics are easy to capture when a user attempts to connect to a network with a web page: operating system and version, IP addresses,

browser, browser fonts, browser plug-ins, time zone, data storage, screen resolution, cookie settings, and HTTP headers. As an example, SecureAuth created SecureAuth Identity Provider (IdP), which provides device authentication. When a user first attempts to access the network with an unknown device, SecureAuth IdP initiates a registration process to capture user data. It records this information and uses it each time the user attempts to log on again with the same device.



TIP Marketers and web analytics services have used similar device fingerprinting services for years. Canvas fingerprinting is a newer method that allows companies to track user activity.

Implementing Access Controls

An access control provides a mechanism to restrict or control access to resources. These resources can be logical, such as files and folders hosted within a network, or physical, such as servers within a locked server room. Access control starts with identification and authentication. Once individuals have been reliability authenticated, you can control their access.

Comparing Subjects and Objects

Two primary terms related to access controls are *subjects* and *objects*. In short, a subject accesses a resource, and an object is the resource being accessed. As a simple example, if a user accesses a project file stored on a server, the user is the subject and the project file is the object.

With this in mind, you would first ensure that subjects are authenticated to validate their identity and then use access controls (such as permissions) on the objects to control the access. The access controls ensure that only certain subjects have access to the objects.

In some examples, an entity can be both a subject and an object. Consider Figure 2-9. A user requesting a web page from a server is clearly a subject. Similarly, the web pages served by the web server are the objects. However, this web server is retrieving data from a back-end database (protected behind a firewall) to build the web page. When the web server retrieves the data, it is a subject and the database is the object.

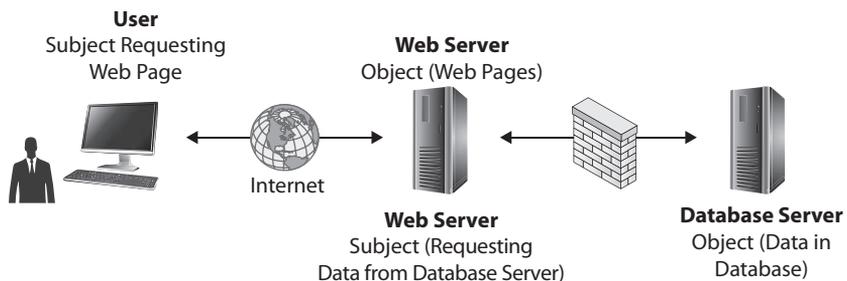


Figure 2-9 Comparing subjects and objects



TIP It may not be apparent, but web servers commonly access back-end databases. For example, Amazon.com hosts web servers to serve the web pages and provides powerful databases to host the data. When you search for a product, the web server sends the query to the database and then formats the reply into the web page. This separation provides security for the database and performance gains for the web server.

Subjects

Access control systems can treat any of the following as subjects:

- Users
- Computers
- Applications
- Networks

You can create accounts to represent most of the subjects. For example, in most networks, each user has at least one user account and each computer has a computer account. Similarly, many applications use service accounts to access resources. You can then assign specific permissions to these accounts for any object they need to access.

Additionally, an access control system can often apply different attributes to the subject to control access further. The following are some subject attributes:

- **Time (temporal)** For example, it's possible to prevent a user from logging on to a system outside of normal working hours.
- **Remote access attribute** Figure 2-10 shows how the Network Access Permission is set to Allow Access for a user in a Microsoft Active Directory domain. This setting allows a user to access a remote access server for dial-in and/or virtual private network (VPN) access.
- **Location** Many systems analyze the user's IP address to determine the user's location. Additionally, some organizations allow users to work from home using a VPN connection, and restrict the users from connecting from other locations. Figure 2-10 shows the Callback Options (in the middle of the dialog box) for a user. It is set to Always Callback to a specific number, which would be the user's home number. When the user calls in, the system will hang up and immediately call the user back at this callback number. If the user isn't at that number, the connection fails.

Many organizations assign administrators two accounts. Administrators log on with a regular user account for day-to-day work, and this account gives them limited access. When they need administrative privileges, they use their administrator account. As an example, Figure 2-11 shows Active Directory Users and Computers (ADUC) in a Microsoft domain. It shows four accounts in the IT Department Organizational Unit. Each of the users has a regular account for day-to-day work and an administrator account used to perform work requiring elevated permissions.

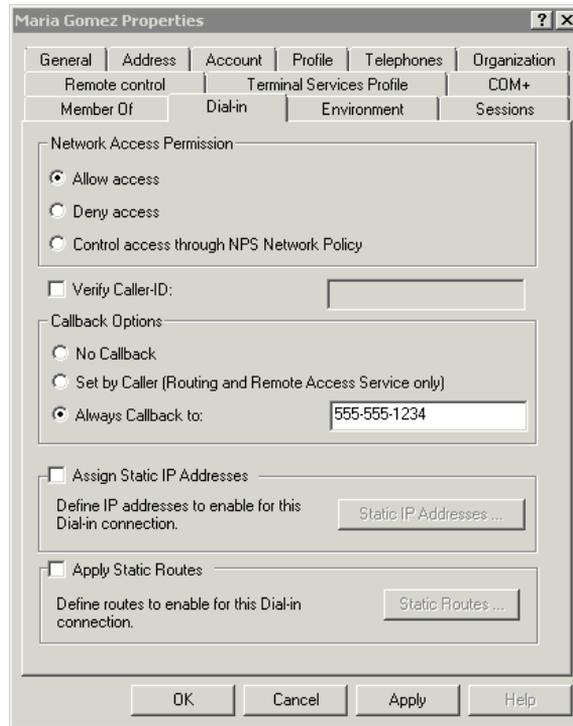


Figure 2-10 Viewing dial-in attributes for a user

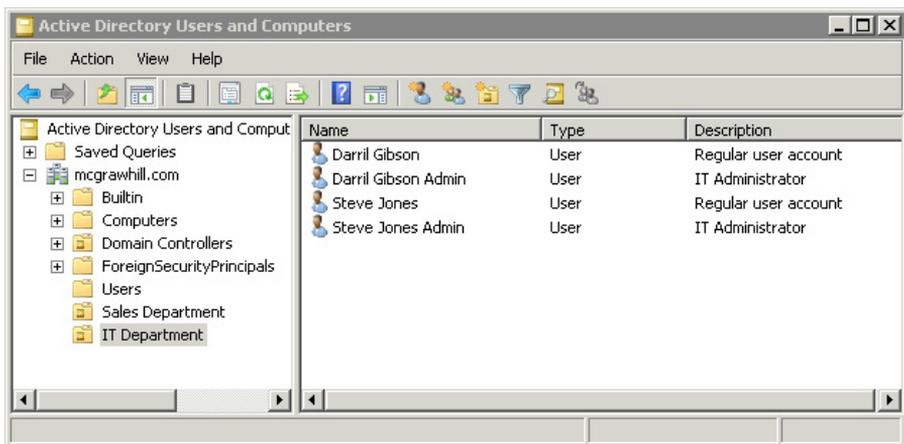


Figure 2-11 Viewing accounts in a domain

Objects

Access control systems can treat any of the following as objects:

- Data (stored in files, folders, and shares)
- Hardware (such as desktop computers, servers, and printers)
- Applications (such as a web server application)
- Networks (such as an Internet connection or an internal connection)
- Facilities (controlled with physical security)

Most objects allow various levels of access. For example, a user might be able to view a web page but not have permission to modify it. Similarly, regular employees can get into the company building but they don't have access to secure server rooms.

It's common to store groups of objects with the same access requirements together. For example, if salespeople need identical access to 100 files, you can store all 100 files in a folder and then assign the permissions to the folder. This simplifies administration.

Logical Access Controls

Logical access controls (also called technical access controls) are implemented with technologies. Many logical access controls use access control lists (ACLs). For example, a router has an ACL with multiple rules that identify the traffic allowed in or out of a network. Similarly, files and folders often use ACLs to identify who can access the resource and what level of access each user has.



EXAM TIP Within IT systems, the security kernel does much of this work. The security kernel is the central part of the operating system that controls access to the system's resources.

Comparing Access Control Models

There are several different access control models available. At this stage of your IT career, you aren't expected to design access based on any of the models. However, you should be aware of the models that exist and have a basic understanding of how they work. The following sections summarize some basic characteristics of these models.

Discretionary Access Control

The Discretionary Access Control (DAC) model provides the most granular level of access control. It is an identity-based model and allows data owners to assign permissions to subjects at the most basic level. For example, you can assign read permission to a single user. File systems such as New Technology File System (NTFS, used by Microsoft) and Network File System (NFS, used on UNIX-based systems such as Solaris and Linux) use the DAC model.



EXAM TIP In the DAC model, users have ownership of the data and can exercise full control over it, including assigning permissions to others. The DAC model allows data owners to assign permissions at the granular level.

DAC uses ACLs, or, more specifically, discretionary access control lists (DACs). As an example, consider Figure 2-12. It shows the permissions assigned to different subjects for the Sales Data object (a folder on a server).



TIP Windows systems identify users and groups with a security identifier (SID) that you'll rarely see. Each entry in the DACL includes the SID and the permission assigned to that SID. However, interfaces such as the one shown in Figure 2-12 perform a lookup to identify the user-friendly name for the SID. They display the name instead of the SID.

A DACL is composed of several access control entries (ACEs), and Table 2-1 shows the underlying ACEs for Figure 2-12. A distinguishing feature of the DAC model is that users have full control over any objects they create. The CREATOR OWNER group is a special group that ensures that users are automatically assigned full control permission to any file or folder they create. Maria Gomez (the sales manager) has full control permission on the folder. The Sales group would include all personnel in the Sales department,

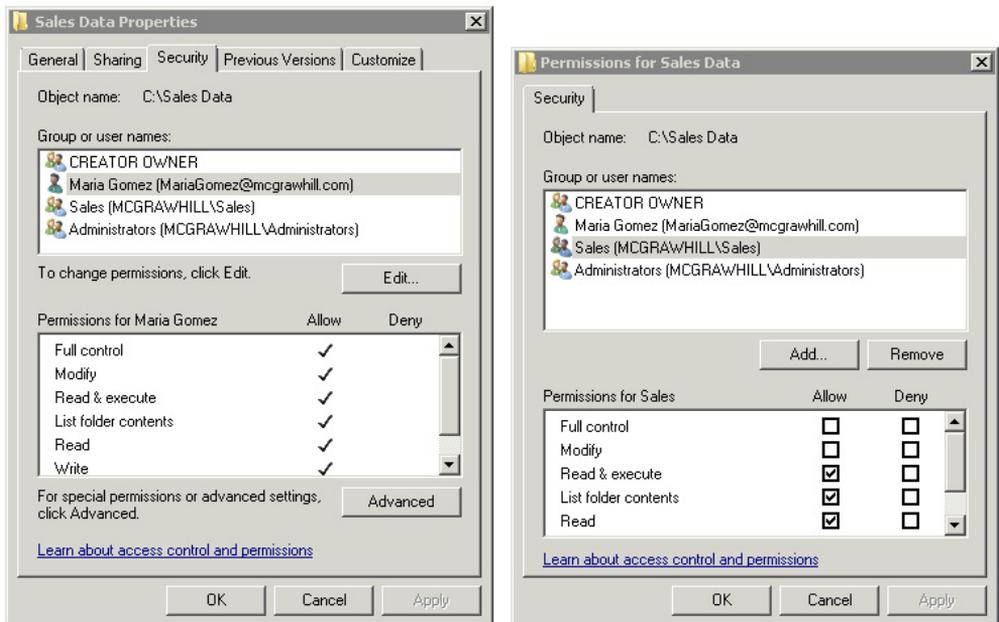


Figure 2-12 Viewing permissions as a DACL

	Subject	Permission(s)
ACE	CREATOR OWNER	Special
ACE	Maria Gomez	Full control
ACE	Sales	Read & execute, list folder contents, read
ACE	Administrators	Full control

Table 2-1 Individual ACEs Creating a DACL

and it has read & execute, list folder contents, and read permissions. The Administrators group has full control permission.



EXAM TIP The DAC model uses a discretionary access control list as shown in Table 2-1. This is a list of subjects along with their rights and/or permissions.

Non-Discretionary Access Control

In non-Discretionary Access Control (non-DAC) models, security administrators control the access granted to users. In contrast, in a DAC model, users have ownership of their resources and users have full control of the resources they own. In general, any model that is not a DAC model is a non-DAC model.

Some operating systems implement non-DAC models for system file access. This prevents malware from taking ownership of any critical or sensitive system files or modifying permissions on any of these files. Users still own and manage their own files using DAC, but the non-DAC model methods protect system files. In other cases, the operating system uses a non-DAC model exclusively, such as Mandatory Access Control (described later in this chapter). Other non-DAC models include role-based, rule-based, and attribute-based models.

Role-based Access Control

The Role-based Access Control (Role-BAC) model uses roles to determine access. Subjects are placed into specific roles and object permissions are granted to the roles. Although the Role-BAC model doesn't provide the granularity offered by DAC, it is easier to implement for large groups of people.



TIP Many operating systems support the use of groups that function as roles. Administrators place user accounts into groups and assign appropriate privileges to the groups.

Consider Figure 2-13. An organization has a group of salespeople in the Sales department. The organization wants each of these users to have access to a printer, a file server, and several folders and files. An administrator adds the accounts for each salesperson to the Sales Role and then assigns appropriate permissions to the Sales Role. These

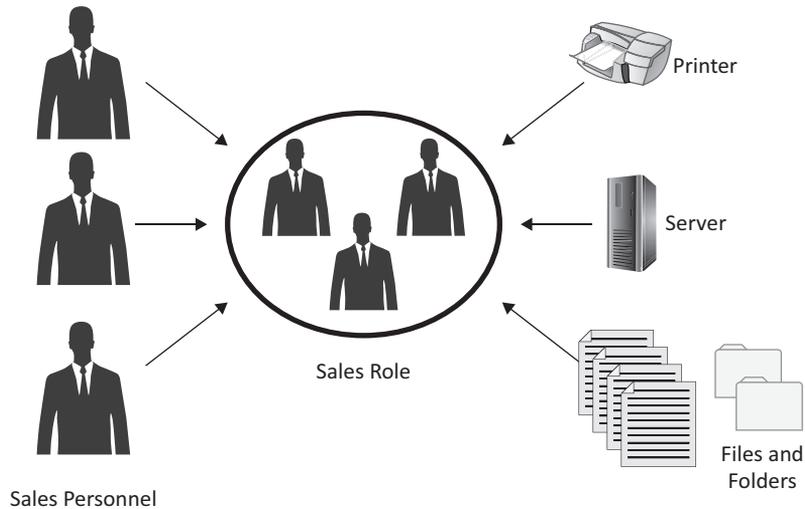


Figure 2-13 The Role-BAC model

permissions apply to all users in the Sales Role without having to assign the permissions individually.

If the organization hires another salesperson, administrators place the new salesperson's user account into the Sales Role. This gives the new employee all the appropriate access in a single step. If a salesperson transfers to a different department, an administrator would remove the account from the Sales Role. When you remove a user from a group or role, it revokes all the permissions they had as a member of the old group or role. An administrator would add the user account to appropriate roles in the new department. This reduces the administrative workload and improves security.



EXAM TIP Role-based models reduce the administrative workload and are very useful in organizations with high employee turnover.

In contrast, if the organization doesn't use roles, but instead forces administrators to assign permissions individually, it becomes very difficult to manage. Looking at Figure 2-13, administrators would have to assign permissions individually for each salesperson at each resource. When a new salesperson is hired, administrators would have to remember all the resources that salespeople need to access and assign permissions individually. When employees transfer, administrators should remove their permissions to resources associated with their previous job. However, without roles or groups, this step is very difficult and often skipped. This results in *permission creep*, where employees just collect more and more permissions. All of this can be challenging with just a few people, but it is unmanageable if your organization has hundreds of users.

Rule-based Access Control

The Rule-based Access Control (Rule-BAC) model is another non-Discretionary Access Control model. Administrators create rules that determine access to resources. As an example, routers have rules within an ACL. These rules identify what traffic the router will pass based on IP addresses, ports, and protocols.



TIP The last rule in a router is an implicit deny rule. It blocks all traffic that isn't explicitly allowed by previous rules. Permissions assigned in DACLS use a similar concept. For example, you explicitly grant permissions to users for a folder. If you don't assign permissions to a specific user, the system blocks that user from accessing the folder.

Attribute-based Access Control

Attribute-based Access Control (ABAC) is a more sophisticated type of Rule-based Access Control. It evaluates subject and object attributes, and grants access based on the value of these attributes. Attributes can be almost any characteristic of a user, the environment, or the resource.

Some systems implement ABAC with plain-language statements to build policies. These policies work as rules, but provide you with a lot more granularity than you can achieve with traditional Rule-based Access Control models. Software-defined networking is becoming very popular and often uses these types of policies. For example, you could create a policy rule that states, "Allow logged-on users to access YouTube videos using smartphones or tablets via the guest network."

Policy statements typically include four elements:

- **Subject** The user accessing a resource. You can use any user property as an attribute. This includes roles, group memberships, management level, and assigned department. The preceding example uses the status of the user—logged on.
- **Object** The resource that the user is trying to access. It can be a file, a database, or an application. The example uses the YouTube application. Notice that this allows the rule to cover much more than just Internet access.
- **Action** What the user is attempting to do. It could be reading or manipulating a file, accessing specific websites, or accessing website applications. The example allows access to the website via the guest network. This could prevent video streaming from overloading the primary network.
- **Environment** Includes everything outside of the subject and object attributes. This is often referred to as the context of the access request. It can include the time, location, protocols, encryption, devices, and communication method. In the example, it specifies smartphones, tablets, and the guest network as environmental attributes.

While the policy statement is stated in plain language, it would typically be stored in a database using an XML format. The eXtensible Access Control Markup Language (XACML) is a specialized XML format used specifically for ABAC.

Mandatory Access Control

The Mandatory Access Control (MAC) model is a non-DAC model that uses labels to identify both subjects and objects. It provides the highest level of security among the models (MAC, DAC, Role-BAC, Rule-BAC, and ABAC) and is commonly used by the U.S. military to ensure that data is protected in mission-critical systems.

The U.S. government uses the following classifications for data, from highest to lowest:

- Top Secret
- Secret
- Confidential
- Unclassified

Just because someone has a Top Secret classification, they don't automatically have access to all Top Secret data. Instead, they are granted access based on their need to know the information for their job. Additionally, it's possible to create sub-classifications or compartments within each classification level.

As an example, Figure 2-14 shows the three classification levels of Top Secret, Secret, and Confidential. Within each level, there are additional labels identifying compartments within the classification level. The Top Secret level includes compartments labeled

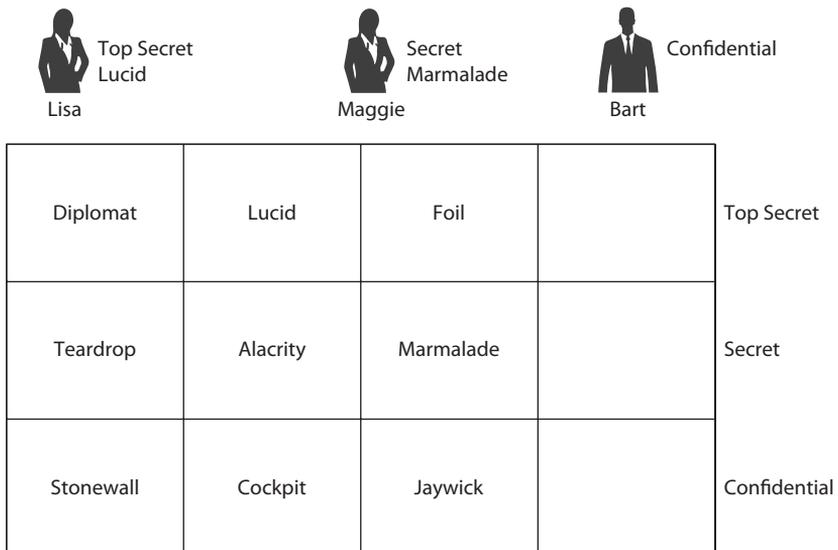


Figure 2-14 Using labels with the MAC model

Diplomat, Lucid, and Foil (names from World War II military operations). The Secret level includes compartments labeled Teardrop, Alacrity, and Marmalade, and the Confidential level includes compartments labeled Stonewall, Cockpit, and Jaywick.

Lisa has both the Top Secret label and the Lucid label. She can be granted access to any data labeled Top Secret and below as long as it doesn't have an additional label. Additionally, she can be granted access to compartmentalized data with the Lucid label. Maggie has the Secret and Marmalade labels. She can be granted access to any data labeled Secret and below as long as it doesn't have an additional label. Additionally, she can be granted access to compartmentalized data with the Marmalade label. Bart has only the Confidential label so he can only be granted access to Confidential data that doesn't have an additional label.



EXAM TIP The MAC model is a non-DAC model that uses labels to control access to data. It is the most secure model when compared to other access control models.

Some of the MAC-based architecture models are Biba, Bell-LaPadula, Clark-Wilson, and the Chinese Wall. Table 2-2 provides an overview of these models, including their primary goals and some of the rules they use. The following sections explain them in a little more depth.

Bell-LaPadula David Elliott Bell and Leonard J. LaPadula designed the Bell-LaPadula model with a primary goal of ensuring confidentiality. It enforces security through two primary rules commonly called *no read up* and *no write down*. Each of these rules compares the subject's clearance with the object's classification.

- **Simple security property rule—no read up** Subjects granted access to any security level may not *read* an object at a higher security level. For example, if Joe is granted Secret access, he cannot read materials classified as Top Secret.
- **The * property (read as “star-property”) rule—no write down** Subjects granted access to any security level may not *write* to any object at a lower security level. For example, if Sally is granted Top Secret access, she cannot create or modify documents classified as Secret.

Model	Goal	Rules
Bell-LaPadula	Confidentiality	No read up, no write down
Biba	Integrity	No read down, no write up
Clark-Wilson	Integrity	Certification (C) rules and enforcement (E) rules
Chinese Wall	Prevent conflict of interest	Access governed by membership in groups to prevent conflicts of interest

Table 2-2 Goals and Rules of Some Access Control Models



EXAM TIP Bell-LaPadula is a MAC model that uses the no read up and no write down rules to enforce confidentiality. These rules compare the subject's clearance with the object's classification. Confidentiality ensures that unauthorized personnel cannot access data.

Biba The Biba model, another MAC-based model, enforces integrity (unlike the Bell-LaPadula model, which enforces confidentiality). Biba includes two rules that are reversed from the Bell-LaPadula model:

- **Simple Integrity Axiom—no read down** Subjects granted access to any security level may not *read* an object at a lower security level, at least not as the authoritative source. For example, a captain of a ship can read orders from an admiral and consider them authoritative and actionable. However, if a seaman recruit tries to issue orders to the captain, the captain will not read them as authoritative.
- **The * Integrity Axiom (read as “star Integrity Axiom”)—no write up** Subjects granted access to any security level may not *write* to any object at a higher security level. For example, a seaman recruit cannot write orders for the captain of the ship. Similarly, the captain cannot write orders for the admiral.



EXAM TIP Biba is a MAC model that uses the no read down and no write up rules to enforce integrity. Integrity protects against unauthorized data modifications.

Clark-Wilson David Clark and David Wilson created the Clark-Wilson model, and its primary goal is information integrity, although it is more stringent than the Biba model. It also helps enforce the separation of duties principle. Chapter 1 described separation of duties. As a reminder, separation of duties is a security principle that ensures that no single person has complete control over a process.

The Clark-Wilson model uses certification rules (identified as C1 through C5) and enforcement rules (identified as E1 through E4) to enforce separation of duties. The certification rules are integrity-monitoring rules, and the enforcement rules are integrity-preserving rules.

These rules are complex, but they work together to ensure there is adequate separation of the elements of any transaction. For example, consider a company that purchases products. The transaction includes someone placing an order, someone receiving the products (and verifying that they were received), and someone paying for the products (after verification that the products were received).

Imagine that Fred has sole responsibility for doing all three tasks in the transaction. This increases the potential for fraud, because he could place an order through a fictitious company, acknowledge receipt (even if no products were received), and deposit the

funds into his own bank account. The Clark-Wilson model ensures that different people perform the separate tasks independently of each other.

Chinese Wall (Brewer-Nash) The Chinese Wall model is documented by Dr. David F. C. Brewer and Dr. Michael J. Nash (and is sometimes referred to as the Brewer-Nash model). Financial services organizations often implement this model to help prevent a conflict of interest, and it helps enforce the separation of duties principle.

Data is classified in the Chinese Wall using different conflict-of-interest classes. If a subject has access to data in one class, he or she cannot access data in a conflicting class.

For example, consider traders in an investment bank who are selling financial securities to investors. These traders should all have access to the same information. In another part of the investment bank, advisors are helping a separate organization with a major merger. If the merger is finalized, it will affect investments sold by the traders. If the traders have access to this information, they could use it to increase investment returns for themselves or their clients. This type of insider trading is illegal, but that isn't enough to stop some people.

A Chinese Wall provides a barrier between these two groups of employees by classifying data. It can be as simple as classifying data as Trader Data and Advisor Data and indicating the two classes have a conflict of interest. Traders have access to Trader Data, and the Chinese Wall blocks their access to the Advisor Data.



EXAM TIP Both the Clark-Wilson model and the Chinese Wall model can be used to enforce the principle of separation of duties. The Chinese Wall model also helps prevent a conflict of interest.

Access Control Matrix vs. Capability Table

Administrators use various tools when implementing access control models. Two common tools are access control matrixes and capability tables. A key difference between the two is that access control matrixes are object based and capability tables are subject based.

An *access control matrix* is a list of objects along with the permissions granted for each object. You can think of an access control matrix as a group of ACLs. Each ACL represents a single object and lists all the permissions for that object.

A *capability table* is a list of subjects, along with the capabilities granted to the subjects. These capabilities include rights and permissions. For example, a capability table can list several groups such as Project Managers, Project Team Leads, and Project Members. It then lists the rights and permissions granted to each of these groups.

Participating in the Identity-Management Life Cycle

Identity management refers to managing accounts through their lifetime. Accounts need to be configured properly when they are first created, managed while in use, and disabled or deleted when they are no longer needed.

Privileges include both rights and permissions. Rights are actions that an account can take on a system, such as backing up files, changing the time, or rebooting the computer. Permissions identify what a user can do with resources, such as reading and writing to a file or printing to a printer.

The principle of least privilege dictates that users should have only the rights and permissions they need to perform their jobs and no more. However, enforcing the principle of least privilege isn't a one-time event. Instead, administrators take steps to enforce this principle throughout the life cycle of any identity.

Identity Proofing

It's important to verify a person's identity before creating a user account. Within an organization, Human Resources (HR) departments typically have processes they use for identity verification. Many HR departments require documents such as a driver's license.

Online banks and retailers often use identity-proofing methods to prevent fraud. One common method to prevent credit card fraud is the use of a credit card verification value (CVV). While criminals might steal databases with credit card numbers, these databases rarely include the CVVs too.

Cognitive passwords (described earlier in this chapter) provide identity proofing after an account has been created. For example, if a user claims to have forgotten a password, the system will challenge the user with cognitive password questions to verify the person's identity.

Provisioning and Authorization

Provisioning refers to creating accounts for users, and *authorization* refers to granting appropriate privileges for the accounts. Many organizations use groups, as described in the "Role-based Access Control" section of this chapter, to streamline the process of granting privileges. For example, administrators can create a group named Sales to represent the Sales department. They add the user accounts for people in the Sales department to the Sales group, and they assign appropriate privileges to the Sales group.

It's common to use template user accounts for different departments. For example, administrators can create an account named SalesTemplate and configure this account with appropriate group membership in the Sales department. When creating an account for a new employee, they make a copy of the template, and it is automatically provisioned with the appropriate group membership.

Some organizations automate account creation. For example, when HR personnel process a new hire, they can enter information about the new employee into an application. This includes basics such as name, address, and the department in which the employee will be working. After the HR person enters the information, the application automatically creates the account and adds the user to the appropriate groups or roles based on the employee's job.

This type of application helps ensure that account-naming conventions are always followed. For example, if the organization wants to create accounts based on a <first name><dot><last name> format (such as darril.gibson), the application will always create it using this convention.

It's important to consider accounts that have elevated permissions, such as administrator accounts. Many organizations require administrators to maintain two accounts. They use one account for day-to-day tasks and have limited privileges, similar to any regular user. They have a second administrator account that they use only when performing administrative work. This helps limit the risks to the administrator account.

For example, if malware infects a system while a user is logged on, the malware can escalate its privileges to the same level as the logged-on user. If the user is logged on with administrative privileges, the malware elevates its privileges to the same level and can then take action as an administrator.

In contrast, if malware infects a system while a regular user is logged on, the malware cannot use the same method to escalate its privileges. If an administrator is logged on as a regular user 90 percent of the time and uses the administrator account only 10 percent of the time, it reduces this risk.

Maintenance and Entitlement

Many organizations have account management policies that dictate how accounts are created and managed through their lifetime. These policies address passwords, account lockouts, and procedures to add and remove privileges.

The “Something You Know” section earlier in this chapter discussed password policies. They ensure that users create strong passwords, change them regularly, and don't reuse the same passwords. It's common to enforce password policies with a technical policy such as with Microsoft's Group Policy. In other words, the password policy starts as a written policy and a technical policy enforces it.

Account Lockout Policies

An account lockout policy locks out an account after too many failed logon attempts. The goal is to prevent attackers from guessing passwords. For example, a policy may dictate that an account is locked out for 30 minutes if someone enters an incorrect password five times in a row.

There are two common settings for account lockout policies:

- **Threshold** This identifies how many incorrect passwords are allowed. For example, if it is set to five, an account is locked out after someone enters the wrong password five times in a row.
- **Duration** This identifies how long the account remains locked out. For example, if the duration is set to 30 minutes, the account will remain locked out for 30 minutes. After 30 minutes, the user will be able to log on with the correct password. In some cases, the duration is configured so that it never expires. The account remains locked out until an administrator unlocks it.

Entitlement and the Principle of Least Privilege

Entitlement refers to the privileges granted to users, and following the principle of least privilege is an important part of entitlement. As mentioned in the “Role-based Access Control” section, using roles (or groups) makes it easier to manage privileges.

Roles and groups provide significant security benefits related to entitlement. For example, suppose that Maria is working in the Sales department. She has appropriate privileges to perform her job there. After earning her SSCP certification, she transfers to the IT department and starts working in IT security. Administrators add her account to IT groups related to her new IT security job and remove her account from groups related to her job in sales.

Removing a user from a group immediately removes all the privileges assigned to the user as a member of that group. In other words, removing Maria from the Sales group removes the privileges that she had while working in the Sales department. The key here is that it's important to remove users from groups when they transfer. If processes aren't in place to monitor entitlement, users end up having more privileges than they need to perform their jobs.

De-provisioning

Accounts de-provisioning refers to disabling and deleting inactive accounts. At a minimum, inactive accounts need to be disabled. When it's determined that the account is no longer needed, it should be deleted.

As an example, if Sally leaves the company to take an executive position in another company, her account needs to be disabled. If it is not, she may be able to access it. If she was able to work from home using a VPN connection, she could still use this VPN connection after she leaves the company. If the account is not disabled and another employee discovers the password, the other employee can log in to the account. Auditing will record any malicious action taken by the other employee as if Sally did it.

Most organizations require that someone disables user accounts as soon as possible after employees leave the company. This includes employees who are terminated and employees who resign. After a specified time, such as 90 days, administrators delete the account. This allows managers to verify the account is not needed prior to deleting it.

In addition to disabling accounts for employees who leave the company, many organizations disable accounts when employees leave for an extended period. For example, if an employee must take an extended leave of absence, the organization disables the account.



EXAM TIP Accounts should be disabled when a user leaves an organization. This includes users who leave an organization amicably or are terminated due to poor performance or other issues.

Participating in Physical Security Operations

While the majority of this chapter focused on logical controls, organizations also use physical access controls to protect assets. Physical access controls include items that you can physically touch. The following list identifies many physical access controls:

- Guards
- Locked doors

- Alarm systems
- Cameras and closed circuit TVs (CCTVs)
- Facilities (controlled with physical security)

Personnel in the organization need to understand and follow appropriate policies related to physical security. As an example, personnel should help prevent tailgating. Tailgating occurs when someone follows someone through a controlled entry point without providing credentials. Imagine that Bob enters the numbers into a cipher lock to open a door. Suzie follows right behind him without entering the numbers into the cipher lock. In this example, Suzie is tailgating.

Some security assessments also address physical security. For example, a security consultant can test physical security by attempting to gain access to secure areas by tailgating or via other methods. If a security consultant is successful, an attacker could also be successful. Chapter 8 covers different types of security assessments.

Chapter Review

Access control systems use identification and authentication. Identification occurs when a subject professes an identity, and authentication occurs when the user proves the identity. Administrators authorize authenticated users access to resources by granting permissions to the resource.

The three factors of authentication are something you know (such as passwords), something you have (such as smart cards), and something you are (using biometrics). Most organizations use written password policies to define requirements. Technical controls implement and enforce password policies, and auditing tools verify compliance with the policies.

Passwords are the weakest authentication mechanism. Static passwords stay the same over a period of time. Dynamic one-time passwords use a different password for each user session and provide much more security than static passwords. Cognitive passwords are based on information that a user knows but an attacker is unlikely to know.

Smart cards, hardware tokens, and proximity cards are examples of something you have. It's common to combine these methods with an additional factor of authentication, providing multifactor authentication. For example, users often enter a PIN along with using a smart card.

Hardware tokens use one-time passwords. A synchronous one-time password is time-based and requires the token and a server to be synchronized. An asynchronous one-time password doesn't require time synchronization but instead uses another method to create the password. Software tokens use protocols such as Time-based One-Time Password (TOTP) and HMAC-based One-Time Password (HOTP). TOTP creates synchronous time-based one-time passwords. HOTP, One-time Password In Everything (OPIE), and S/KEY create asynchronous one-time passwords.

Biometrics methods measure physical characteristics of individuals, such as their fingerprints, palms, retinas, and irises. These characteristics are in the something you are factor of authentication. When evaluating the effectiveness of biometrics, you should consider

the type 1 and type 2 errors. The False Rejection Rate (FRR, or type 1 error) refers to the percentage of times a system falsely rejects a known user. The False Acceptance Rate (FAR, or type 2 error) refers to the percentage of times a biometric system falsely identifies an unknown user as a known user. The Crossover Error Rate (CER) indicates the point where the FAR and FRR are equal. Lower CERs indicate a better biometric system.

Single sign-on (SSO) allows a user to authenticate once for a system. The user can then access any resources in the system (as long as the user has access) without authenticating again. Kerberos is a popular SSO protocol used within networks. A Kerberos server issues symmetric encryption keys as a Key Distribution Center (KDC) and issues tickets as a ticket-granting server. Kerberos provides authentication on a network and contributes to the confidentiality and integrity of information.

Many SSO systems use federated access, providing centralized authentication for different systems. Security Assertion Markup Language (SAML) uses an XML-based data format to share SSO information between servers within a federated identity management system. While SAML is primarily used for authentication, it can also exchange authorization information between different entities.

Access controls enforce security. For example, a security kernel (a central part of an operating system) enforces security for the operating system by monitoring subjects and objects. Examples of subjects are users, computers, and applications. Examples of objects include data, hardware, and facilities. Access controls can be logical (implemented with technology such as a security kernel) or physical (such as locked doors).

Several access control models are in use, including discretionary (DAC), non-discretionary (non-DAC), mandatory (MAC), role-based (Role-BAC), rule-based (Rule-BAC), and attribute-based (ABAC) models. The DAC model provides the most granular control and is an identity-based model. Individual users own the objects (such as files), and as the data owners, they can assign permissions to subjects (such as other users) as desired. DAC is used with file systems such as NTFS with Microsoft systems and NFS with Linux systems.

The RBAC model uses roles. Subjects (such as users) are placed into roles, and permissions to objects are assigned directly to the role. In non-DAC models, security administrators control access. An ABAC model is a sophisticated type of Rule-based Access Control model. Policies identify attributes of subjects, objects, and the environment to control access.

The MAC model provides the highest level of security when compared to DAC and RBAC models. Both subjects and objects are assigned labels, and when the labels match, access is granted. Several architectures are based on the MAC model. The Bell-LaPadula model has a primary goal of ensuring confidentiality and uses rules of no read up and no write down. The Biba model has a primary goal of ensuring integrity and uses rules of no read down and no write up. The Clark-Wilson model provides integrity by using certification and enforcement rules to enforce integrity and separation of duties. The Chinese Wall helps prevent conflicts of interest by preventing access to data organized in conflict-of-interest classes.

Identity management includes provisioning, maintenance, and entitlement. Provisioning includes creating accounts and providing appropriate access. Many organizations use roles

or groups to manage access to resources. Account management ensures that accounts use appropriate password policies and that accounts are disabled when the employee is no longer working for the company. Entitlement refers to privileges granted to users and helps ensure that the principle of least privilege is followed.

Questions

1. A user professes an identity by entering a user logon name and then enters a password. What is the purpose of the logon name?
 - A. Authentication
 - B. Accountability
 - C. Identification
 - D. Accounting
2. Access controls protect assets such as files by preventing unauthorized access. What must occur before a system can implement access controls to restrict access to these types of assets?
 - A. Identification and authentication
 - B. Identification and accountability
 - C. Authentication and accounting
 - D. Accountability and availability
3. Users are required to enter a different password each time they log on. What type of password is this?
 - A. Static password
 - B. Cognitive password
 - C. Passphrase
 - D. Dynamic password
4. Authentication includes three types, or factors. Which of the following best describes these authentication methods?
 - A. Something you say, something you think, and something you are
 - B. Something you know, something you have, and something you type
 - C. Something you know, something you say, and something you are
 - D. Something you know, something you have, and something you are
5. Which of the following choices does NOT ensure that a password is strong?
 - A. Ensuring that the password is of a sufficient length
 - B. Ensuring that the password is changed frequently
 - C. Ensuring that the password has a mixture of different character types
 - D. Ensuring that the password does not include any part of the user's name

6. What can be used to prevent a user from reusing the same password?
 - A. Minimum password age
 - B. Maximum password age
 - C. Password length
 - D. Password history
7. What form(s) of authentication are individuals using when they authenticate with a hardware token and a password?
 - A. Something they have only
 - B. Something they know only
 - C. Something they have and something they know
 - D. Something they have and something they are
8. An organization uses a biometric system with a one-to-many search method. What does this system provide for the organization?
 - A. Authentication
 - B. Accountability
 - C. Authorization
 - D. Identification
9. An organization has been using an iris scanner for authentication but has noticed a significant number of errors. Assuming the iris scanner is a high-quality scanner, which of the following could affect its accuracy?
 - A. False Acceptance Rate (FAR)
 - B. False Rejection Rate (FRR)
 - C. Sunlight shining into the scanner
 - D. Faulty laser beam
10. Which of the following metrics identifies the number of valid users that a biometric authentication system falsely rejects?
 - A. FAR
 - B. FRR
 - C. CER
 - D. AAA
11. Which of the following biometric methods has the lowest CER?
 - A. Iris scan
 - B. Handwriting analysis
 - C. Keystroke dynamics
 - D. Thumbprint scan

12. What is SSO?
 - A. A system that requires user credentials once and uses the same credentials for the entire session
 - B. An authentication system that requires users to use different credentials for each resource they access
 - C. A secure system used for operations
 - D. Any network that employs secure access controls
13. What type of service does Kerberos provide?
 - A. Authentication
 - B. Accounting
 - C. Availability
 - D. Accountability
14. Of the following choices, what most accurately identifies the major drawback of SSO systems?
 - A. It allows users to access multiple systems after logging on once.
 - B. It increases the difficulty for users to log on.
 - C. It increases the administrative workload.
 - D. It risks maximum unauthorized access with compromised accounts.
15. What type of access control is identity based?
 - A. Discretionary
 - B. Non-discretionary
 - C. ABAC
 - D. Biba
16. What is the primary goal of the Bell-LaPadula model?
 - A. Enforce separation of duties
 - B. Enforce two-factor authentication
 - C. Enforce confidentiality
 - D. Enforce integrity
17. Which of the following models helps enforce the principle of separation of duties?
 - A. Chinese Wall and Clark-Wilson
 - B. Chinese Wall and Biba
 - C. Clark-Wilson and Bell-LaPadula
 - D. Biba and Bell-LaPadula

18. Which of the following statements is true?
- A. An access control matrix is object based and a capability table is object based.
 - B. An access control matrix is subject based and a capability table is object based.
 - C. An access control matrix is object based and a capability table is subject based.
 - D. An access control matrix is subject based and a capability table is subject based.
19. Which of the following will disable an account if an attacker tries to guess the password multiple times?
- A. A password policy
 - B. An account lockout policy
 - C. A password history
 - D. De-provisioning accounts
20. Which of the following actions is most appropriate if an employee leaves the company?
- A. Delete the user's account as soon as possible.
 - B. Disable the user's account as soon as possible.
 - C. Change the user's password as soon as possible.
 - D. Change the user's permissions as soon as possible.

Answers

1. **C.** The logon name provides identification of the user. When combined with the username, the password provides authentication. Accountability is possible if a system can identify users and track their activities. Accounting is provided by logging after proper identification and authentication.
2. **A.** Identification and authentication must occur before a system can implement access controls. Identification is the act of a user professing an identity, and authentication occurs when an authentication system verifies the user's credentials (such as a username and password). Without proper identification and authentication, it isn't possible to restrict access to specific subjects. Accountability is not possible if an authentication system does not identify and authenticate users. Similarly, you can't provide accurate accounting if users haven't been identified and authenticated.
3. **D.** A dynamic password is a one-time password that changes for each session. A static password stays the same over a period of time, such as 30 days. A cognitive password includes information that a user would know but isn't necessarily public knowledge. A passphrase is a static password using a long string of characters that has meaning to the user.

4. **D.** The three factors of authentication are something you know, something you have, and something you are. Something you think, something you type, or something you say are not authentication factors.
5. **B.** A password should be changed regularly, but doing so doesn't ensure the password is strong. For example, if a user changes a password from "1234" to "4321," it is not strong. The other options all contribute to the strength of a password.
6. **D.** Password history remembers users' previous passwords and prevents them from reusing passwords. The minimum password age is used with the password history to prevent users from changing their password repeatedly to get back to the original password. It is often set to one day. The maximum password age identifies when users must change their passwords. The password length identifies the minimum number of characters in the password.
7. **C.** The two factors of authentication are something they have (the hardware token) and something they know (the password). The third factor of authentication is something you are (using biometrics), but neither a hardware token nor a PIN uses biometrics.
8. **D.** A biometric system used for identification uses a one-to-many search method. Biometric systems used for authentication use a one-to-one search method. Once a system identifies and authenticates a user, biometric systems are not used for accountability or authorization.
9. **C.** Lighting affects the accuracy of an iris scanner, so sunlight shining into the scanner's aperture will affect the accuracy. The FAR and FRR indicate specific types of errors. An iris scanner does not use a laser beam.
10. **B.** The False Rejection Rate (FRR, also called a type 1 error) refers to the percentage of times a biometric system falsely rejects a known user. The False Acceptance Rate (FAR, also called a type 2 error) refers to the percentage of times a biometric system falsely identifies an unknown user. The Crossover Error Rate (CER) identifies where the FAR matches the FRR. AAA refers to the three A's of security—authentication, authorization, and accounting.
11. **A.** Iris scans are the most accurate of the items listed and have the lowest Crossover Error Rate (CER). Behavioral biometrics methods such as handwriting analysis and keystroke dynamics are not as accurate as biometric methods that use physical characteristics. Fingerprint scans and thumbprint scans are commonly used but are not as accurate as iris scans.
12. **A.** Single sign-on (SSO) requires users to log on once, and it uses the same credentials for any other resources accessed during the session. Users are not required to use different credentials for each resource with SSO. SSO is not a network.
13. **A.** Kerberos provides authentication. Accounting and accountability are possible if a system can identify users and track their activities, but Kerberos doesn't provide logging of user activities. Availability ensures systems are operational when needed. While not listed as one of the answers, Kerberos also contributes to confidentiality and integrity.

14. **D.** A major concern with SSO systems is that if any single account is compromised, it maximizes the potential unauthorized access. A primary benefit is that SSO allows users to access multiple systems after logging on once. It decreases the difficulty for users to log on and decreases the administrative workload.
15. **A.** A Discretionary Access Control (DAC) model assigns permissions to identities, making it an identity-based model. Not all non-discretionary models assign permissions to identities, so B is not the best answer. Attribute-based Access Control (ABAC) focuses on attributes of subjects, objects, and/or the environment. Biba is a MAC-based model that uses labels assigned to both subjects and objects to enforce integrity.
16. **C.** The Bell-LaPadula model has a primary goal of ensuring confidentiality. The Clark-Wilson and Chinese Wall access control models help enforce the principle of separation of duties. The Biba model enforces confidentiality. Access control models do not handle authentication.
17. **A.** Both the Clark-Wilson model and the Chinese Wall model enforce the principle of separation of duties. The Clark-Wilson model also enforces integrity, and the Chinese Wall model also helps prevent conflicts of interest. Biba enforces integrity. Bell-LaPadula enforces confidentiality.
18. **C.** An access control matrix is object based and a capability table is subject based. An access control matrix is a list of objects along with the permissions granted for each object. A capability table is a list of subjects along with the capabilities granted to the subjects.
19. **B.** An account lockout policy can disable an account if an attacker (or a user) enters the wrong password too many times. The threshold is often set to three or five, causing an account to be locked out after a user enters the wrong password three or five times, respectively. A password policy ensures that users create strong passwords and regularly change their password. Password history prevents users from reusing the same password. De-provisioning refers to disabling and deleting inactive accounts.
20. **B.** User accounts should be disabled as soon as possible after the user leaves the company under any circumstances. The account should not be deleted until it's determined to be no longer needed. Changing the password without disabling the account stills allow the account to be used. Disabling the account removes the access and is more direct than changing permissions.

Basic Networking and Communications

In this chapter, you will learn about

- The OSI and TCP/IP models
- Bus, star, tree, token ring, and mesh network configurations
- Common protocols in the TCP/IP suite
- Ports and protocol numbers used with common protocols
- Different network and trust architectures such as the Internet, DMZs, and transitive trusts
- Wireless security protocols such as WPA, WPA2/802.11i, AES, and TKIP
- Wireless technologies such as Bluetooth, RFID, 802.11, WiMAX, GSM, 3G, and NFC

The OSI Model

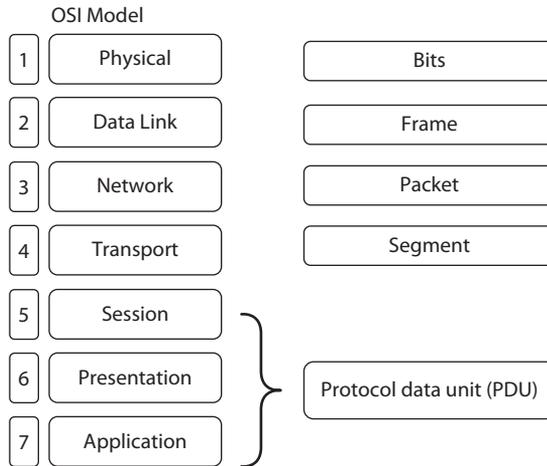
The International Organization for Standardization (ISO) developed the Open Systems Interconnection (OSI) Model as a standard model for network communications. It describes a framework for connecting multiple computers to each other and allows dissimilar networks to communicate with each other. It includes seven layers, and each layer focuses on a set of specific tasks or functions.

Data is packaged, or encapsulated, at each layer with additional data that the layer requires. Generically, the encapsulated data is a *datagram*. However, datagrams have specific names at each layer. For example, at the lowest layer (the Physical layer), data arrives as *bits* (ones and zeros). The next layer (the Data Link layer) adds more data, and packages the data into *frames*. Figure 3-1 shows a diagram of the OSI Model identifying the layers and the data encapsulation names. The following subsections describe these different layers.



TIP Many people memorize the layers of the OSI Model with a mnemonic. For example, All People Seem To Need Data Processing (for Application, Presentation, Session, Transport, Network, Data Link, and Physical) is one. Please Do Not Throw Sausage Pizza Away (for Physical, Data Link, Network, Transport, Session, Presentation, and Application) is another.

Figure 3-1
OSI Model



For the SSCP exam, you don't need to know the OSI Model as in-depth as you would for a networking certification such as Cisco's CCNA or CompTIA's Network+. However, you should have a basic understanding of the model, including what occurs at each layer. This section introduces each of the OSI Model layers. The "Reviewing Basic Protocols and Ports" section later in this chapter includes a figure that maps protocols to specific layers.

The Physical Layer (Layer 1)

The Physical layer provides a direct connection to the network. It defines how data is passed onto the network, including the transmission method, such as via wired or wireless techniques. The Physical layer includes physical items such as cabling, hubs, repeaters, and wireless radio waves required to provide a connection to the network.

Security at this layer focuses on protecting access to the network connections. If attackers are able to connect a protocol analyzer or sniffer via unprotected cabling, they can intercept data going across the network.



TIP A sniffer can capture and analyze packets sent over both wired and wireless networks. For example, if systems send credentials such as a username and password over a network in clear, readable text, the attacker can learn the credentials by capturing and analyzing the traffic. In contrast, attackers are unable to read the data if systems encrypt it prior to transmission, making the text unreadable.

The following list describes the different types of media used for data transmissions:

- **Coaxial cable** Coaxial cable is similar to the cable used for cable television. It's an older technology that is rarely used on networks today, though it was commonly used in early bus networks (described later in this chapter). It is relatively easy for an attacker to tap into a coaxial cable connection.

- **Twisted pair** This is the most commonly used media for data transmissions. Unshielded twisted pair (UTP) does not have shielding and is susceptible to electromagnetic interference (EMI) and radio frequency interference (RFI). Shielded twisted pair (STP) provides protection against EMI and RFI. The EIA/TIA-568-B standards define several categories of twisted pair cables. Common standards are Cat 5e and Cat 6. It is relatively easy for an attacker to tap into a twisted pair connection.
- **Fiber-optic cabling** Fiber provides the highest bandwidth and can support cable runs up to 40 kilometers. Although more expensive than copper cable such as twisted pair, fiber is becoming more and more popular. It is free from EMI and RFI issues and it is harder to tap into than wireless, twisted pair, and coaxial.
- **Wireless transmissions** Wireless transmissions use radio frequency waves to transmit the signals over the air. Because they travel over the air, they are the easiest to intercept.



EXAM TIP Cable standards are defined on the Physical layer. Wireless transmissions are the easiest to intercept, and fiber connections are the most difficult to tap into.

Data traveling on the Physical layer is packaged and transmitted as bits (ones and zeros) and bit streams. The Physical layer converts the ones and zeros from the media and passes a bit stream to the Data Link layer, which packages the data as a frame.

The Data Link Layer (Layer 2)

The Data Link layer provides reliable delivery of data across the network and establishes communication links between devices over a network. It includes two sublayers. The media access control (MAC) sublayer interacts with the Physical layer and includes the MAC address (also called the physical address or hardware address). The Logical Link Control sublayer interacts with the network layer.

Although manufacturers assign a MAC address to network interface cards (NICs), it's worth noting that it's possible to assign a different MAC address to a NIC by using software. This allows an attacker to spoof the source MAC address so that traffic appears as though it's coming from a different computer.

Basic layer 2 switches work at this layer. As traffic comes into a switch on a physical port, the switch identifies the MAC address of the computer sending the data. It then creates a table mapping the MAC address to the physical port. When the switch receives unicast traffic for that computer, the switch sends the data to that port only.

Data traveling on the Data Link layer is packaged and transmitted as frames. The Data Link layer includes orderly delivery of frames and error notification if frames are not delivered properly.



EXAM TIP The Data Link layer is the only layer that includes sublayers. It includes the media access control (MAC) sublayer, which defines MAC addresses (or physical addresses), and the Logical Link Control sublayer. Data is packaged in frames, and the Data Link layer ensures orderly delivery of frames.

The Network Layer (Layer 3)

The Network layer provides routing for data across a network by analyzing the Internet Protocol (IP) address and determining the best route to the destination computer. This layer also controls the flow of data across the network.

Routers and layer 3 switches are the primary hardware devices used at this layer. Routers use access control lists (ACLs) for basic packet filtering to control traffic. An ACL can block or allow traffic based on IP addresses, logical ports, and even some protocols such as Internet Protocol security (IPsec, identified with protocol numbers 50 and 51). Data traveling on the Network layer is packaged and transmitted as packets.

IP is the protocol used for addressing. Currently both IPv4 and IPv6 addresses are in use on the Internet and in many internal networks.



EXAM TIP Routers and layer 3 switches operate on the Network layer and determine the best route based on the destination IP address (also known as the logical address). Both IPv4 and IPv6 addresses are defined on this layer.

The Transport Layer (Layer 4)

The Transport layer provides reliable end-to-end communication services for applications, and includes error detection and recovery mechanisms. Data traveling on the Transport layer is packaged and transmitted as segments. Many of the logical port assignments for upper-layer protocols occur on this layer. Two primary protocols used at this layer are Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).



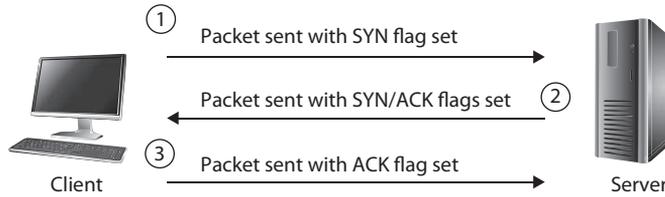
TIP Well-known ports are logical numbers matched to specific protocols. For example, HyperText Transfer Protocol (HTTP) uses the well-known port of 80. By blocking port 80, the router's ACL blocks HTTP traffic that is using port 80. The port number is embedded in the datagram.

TCP

TCP is a connection-oriented protocol that provides guaranteed, reliable communication for devices on a network. It uses packet sequencing, and the destination acknowledges every packet that it receives. A three-way handshake establishes a TCP connection.

In Figure 3-2, a client is establishing a TCP session with a server. In the first step of the three-way handshake, the client sends a packet with the synchronize (SYN) flag set. The server responds with both the acknowledge (ACK) and SYN flags set. Last, the client sends back a packet with the ACK flag set. At this point, the session is established.

Figure 3-2
TCP three-way
handshake



EXAM TIP In a common denial of service (DoS) attack known as a SYN flood attack, the client withholds the third packet. The server is left with a half-open session consuming system resources while it waits for the third packet. However, instead of sending the third packet, the attacker floods the system with more SYN packets. Chapter 5 covers DoS attacks in more depth.

If another protocol requires guaranteed delivery, it uses TCP. For example, IP doesn't include the ability to ensure guaranteed delivery. Instead, IP uses TCP to provide the guaranteed delivery. If the protocol does not require guaranteed delivery, it uses UDP.

UDP

UDP is a connectionless protocol. Instead of checking to see whether a connection exists with another system before sending data, UDP simply sends it. Another way of thinking of this is that UDP gives a "best effort" to send the data, but it does not verify delivery. This could be like your sending a text or an e-mail to a friend to pass on some information. There's a good chance your friend will receive the message, but you don't have any automatic verification.

Many diagnostic protocols use UDP because they don't need to establish a session when performing diagnostic tests. You'll also see UDP used in simpler protocols such as Trivial File Transfer Protocol (TFTP), a simple, streamlined version of File Transfer Protocol (FTP). TFTP sends small files such as configuration files to network devices. FTP uses TCP, but TFTP uses UDP. UDP is also common with streaming data such as streaming audio or video. With streaming data, it's acceptable to lose some bits occasionally here and there. For example, a streaming video file may hiccup occasionally when some bits are lost, but the overall file still plays.



TIP The terms *connection-oriented* and *connectionless* have nothing to do with a wired or wireless connection. Instead, they refer to the logical nature of a connection. TCP ensures both sides can communicate before starting a session. UDP simply sends the data without checking to see whether a logical connection exists.

The Session Layer (Layer 5)

The Session layer establishes and maintains sessions between applications or software processes on the local and remote systems. It keeps the session information for one application

separate from the session information from other applications. Remote Procedure Call (RPC) is one of the few protocols that operates directly on this layer. Applications can use RPC to request a service from an application running on a different computer.

Data traveling on the application, presentation, and session layers is packaged and named protocol data units (PDUs).

The Presentation Layer (Layer 6)

The Presentation layer standardizes data presentation for the application layer. It translates data using standards such as American Standard Code for Information Interchange (ASCII), Extended Binary Coded Decimal Interchange Code (EBCDIC), Joint Photographic Experts Group (JPEG), Moving Picture Experts Group (MPEG), and so on. Some data encryption and decryption occurs at this layer. Additionally, data can be compressed and decompressed at this layer for better network performance.

The Application Layer (Layer 7)

Network applications use the Application layer to implement specific user applications. This layer provides the services needed for applications that communicate across a network.

The Application layer supports many additional security features. These include authentication, access control with mechanisms such as permissions, encryption for confidentiality, hashing for integrity, and the use of digital signatures for nonrepudiation.



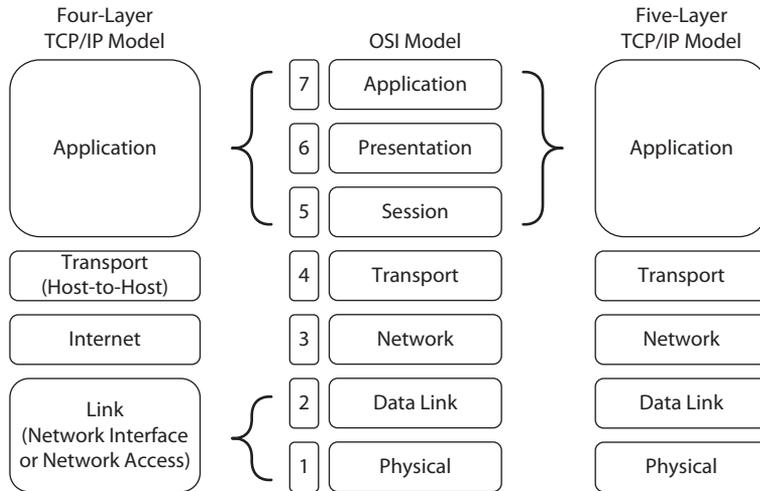
TIP Chapter 14 covers many cryptographic topics, including encryption, hashing, and digital signatures.

Comparing the OSI and TCP/IP Models

The Defense Advanced Research Projects Agency (DARPA) created the TCP/IP Model. It has a similar goal as the OSI Model—to provide a common framework for connecting multiple computers together. Request for Comments (RFC) 1122 and RFC 1123 identify the TCP/IP Model with four layers that map to the OSI layers. Some documentation identifies the TCP/IP Model with five layers as shown in Figure 3-3, but the RFCs showing four layers are authoritative.

Looking at the four-layer TCP/IP Model, you can see that the upper-three layers of the OSI Model map directly to the Application layer of the TCP/IP Model. The OSI Transport layer maps to the TCP/IP Transport layer (also known as the Host or Host-to-Host layer). The OSI Network layer maps to the TCP/IP Internet layer. Finally, the OSI Data Link and Physical layers map to the TCP/IP Link layer (also known as the Network Interface or Network Access layer). Documentation that shows five layers typically names the bottom-four layers the same as they are on the OSI Model. In essence, they perform the same tasks.

Figure 3-3
TCP/IP Model
compared to the
OSI Model



EXAM TIP Although RFCs 1122 and 1123 identify the TCP/IP Model with four layers, some references listed in the Candidate Information Bulletin (CIB) identify it with five layers. The TCP/IP Model is also called the TCP Model and the DoD Model. DoD is an acronym for the Department of Defense.

Network Topologies

The network topology refers to the layout of a network. It includes items you can touch such as cables and network devices. The primary configuration you'll see in most networks today is star, but you may also see bus, tree, token ring, and mesh configurations. In addition to understanding the topologies, it's important to understand how Ethernet fits into the mix.

Ethernet

The IEEE 802.3 family of standards defines Ethernet. It is the most widely used standard for local area networks (LANs). A working group drafted the first standard in 1973, and the Institute of Electrical and Electronics Engineers (IEEE) has steadily released updates since then. For example, Gigabit Ethernet over twisted pair (802.3bp) and 10 Gigabit Ethernet over fiber (802.3ae) define the standards used for sending data over a network with progressively higher bandwidths.



TIP Modifications to the 802.3 standard are identified with lowercase letters such as 802.3a, 802.3b, and so on. While the standard alphabet has only 26 letters, there have been more than 26 modifications to the 802.3 standard. The modification after 802.3z is 802.3aa. Similarly, the modification after 802.3az is 802.3ba.

Topology vs. Topography

Network concepts such as bus, star, and token ring were included in previous versions of the SSCP domain objectives as both topologies and topographies. The 2012 domain objectives only listed topographies. The 2015 domain objectives include both topologies and topographies.

Two questions came to mind for me: 1) What's the difference between topology and topography? 2) Does the change imply a different focus on these concepts for a network?

In the context of networking, *topology* refers to the arrangement of elements in a network. The logical topology identifies how data flows through the network (such as over a single-line bus or in a ring). The physical topology identifies the actual placement of devices in the network.

Topography usually refers to maps that show physical characteristics of a land mass, such as the elevations of mountains or the physical features of an island. It can also refer to a detailed description of map features, but you'll find it very difficult to find any documentation referring to a network topography.

I've concluded that some people are just merging these terms. While some people might want to debate the correctness of the terms, the good news is that when you're taking the SSCP exam, it doesn't matter. Exam questions might use the term topology or topography, but what you really need to know is the difference between concepts such as bus, star, tree, token ring, and mesh.

When discussing Ethernet, it's important to understand Carrier Sense Multiple Access (CSMA) and collisions. A collision occurs when two devices try to communicate at the same time over the same medium. Because both devices are sending traffic, the data from both systems collides with each other and none of the data is readable. There are two ways of dealing with collisions if the systems are using the same wire for sending and receiving: detecting them or avoiding them.

CSMA/CD

Ethernet networks support Carrier Sense Multiple Access/Collision Detection (CSMA/CD). In CSMA/CD, each of the devices on the network can detect collisions. As soon as a device detects a collision, it sends a signal on the wire letting all other systems know that a collision occurred. Each device then waits a random amount of time before resending.

However, most switches and computers today use full-duplex connections. In other words, the cable connecting a switch and a computer includes two sets of wires, with one set of wires used for receiving and another set of wires used for transmitting. If full duplex is used, and each device is connected directly with only one other device, collisions don't occur on this cable.

Unfortunately, some older devices still use half duplex. In a half-duplex connection, a single set of wires is used for both transmitting and receiving. It's possible for one system to transmit at the same time as another device is transmitting, resulting in a collision. Ethernet supports CSMA/CD for backward compatibility.

CSMA/CA

In Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA), systems listen before transmitting data. If they sense that other systems are transmitting data, they wait a random amount of time to check again before sending traffic. Wireless networks such as 802.11 networks use CSMA/CA.

Wireless networks also support Request To Send (RTS) and Clear To Send (CTS) to negotiate traffic. When a computer wants to transmit data, it sends out the RTS signal. The destination replies with a CTS signal and waits for the transmission.

Bus

In a bus configuration, all of the hosts are connected together via a single shared cable connection as shown in Figure 3-4. Earlier networks commonly used the bus configuration because it was easy to set up. However, it is very difficult to troubleshoot, so you won't see it used to connect entire networks together today.

The cable used in a bus configuration is typically coaxial cable. This is similar to a cable used for cable television. T connectors connect each computer on the network. The cable connects to one end of the T, the base of the T connects to the computer, and the other end of the T goes to the next computer.

Notice that each end of the bus includes a terminator. When an electrical signal travels down a wire, it will reflect back when it reaches the end unless the wire has a terminator. This reflected signal corrupts other data transmissions on the bus. The terminator is a connector that prevents this signal reflection. Removing the terminator effectively takes the entire bus network down.

Consider how the computers are connected with a T connector to the main bus. If any one of these T connections is disconnected, you no longer have a single bus with two terminators. Instead, you have two buses, and each one only has a single terminator instead of two terminators. In other words, none of the computers will be able to communicate on the network.

If you only have three or four computers on your network, you can probably locate the problem relatively easily. However, if the network includes several hundred computers, it can take quite a long time to discover the source of the problem.

Older Ethernet standards for bus configurations include 10Base2 (ThinNet) and 10Base5 (ThickNet). These use thin and thick coaxial cables, respectively, for bus configurations.



EXAM TIP Bus networks are expensive to maintain because of how difficult it is to troubleshoot problems. A single break in the cable takes down the entire network, reducing availability.

Figure 3-4
Bus configuration

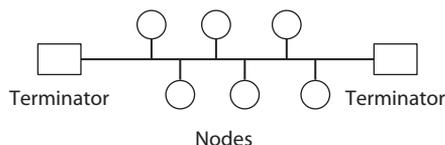
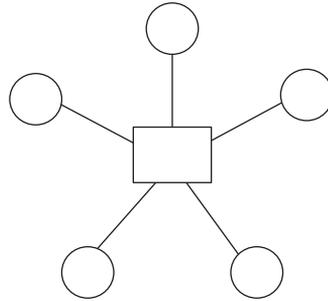


Figure 3-5
Star configuration



Star

In a star configuration, all the devices are connected to a central device such as a hub or switch, as shown in Figure 3-5. When drawn this way, it is reminiscent of a star, which is how it gets its name.



NOTE Many network diagrams are drawn similar to a bus configuration diagram but are actually configured in a star. In other words, network diagrams don't indicate whether the network is configured as a bus, star, or token ring. Instead, they show the relevant components, such as the routers and networks.

A star configuration is much more common in typical computer networks today. Additionally, the devices typically connect with UTP cable or STP cable. UTP and STP cable use RJ-45 connectors. Common twisted pair cable categories are Cat 5e, 6, and 6a. Cat 5e cable supports speeds up to 1 gigabit and Cat 6 and 6a cables support speeds up to 10 gigabit. The maximum length of these cables is limited to 100 meters (about 330 feet).

A typical star configuration uses a switch as the central component, and computers connect to the switch as nodes in the star with UTP or STP cables. Each of the computers communicates with other computers on the network via the switch. Network devices include components to terminate connections automatically, so external terminators are not required in a star configuration as they are in a bus configuration.

Networks configured in a star are much easier to troubleshoot than a bus network. If someone disconnects a cable from a single computer, only that computer is affected. All the other computers still work. However, compared to the bus configuration, you need much more cable. Each device uses a separate cable from it to the device, which is typically stored in a secure location such as a locked wiring closet or a server room.

The central device in a star configuration could be a hub, but most organizations have replaced hubs with switches. When a hub is used, all traffic that goes in one port goes out all the other ports of the hub. An attacker could install a protocol analyzer (sniffer) on any computer connected to the hub and capture all the data going through the hub. However, when organizations replace hubs with switches, the only unicast traffic that will reach the attacker's computer is traffic that is directly addressed to or from the attacker's computer.

This improves network efficiency by limiting traffic to computers, and it reduces the potential success of sniffing attacks.

It is possible to configure most switches with a mirrored port. This sends a copy of all traffic through the switch to the mirrored port for monitoring. If an attacker has physical access to the switch, the attacker can connect to the monitoring or mirrored port and sniff all the traffic. Because of this, organizations keep switches in secure areas, such as in a locked wiring closet or a locked server room. Securing a switch in a wiring closet or server room provides physical security.

Ethernet standards for star configurations use twisted pair cable such as 10Base-T, 100Base-TX, and 1000Base-T, supporting bandwidths of 10 Mbps, 100 Mbps, and 1 Gbps, respectively.

Tree

A tree topology is a combination of a bus topology and a star topology. Instead of connecting multiple computer nodes in the bus configuration, it connects multiple star networks along a type of a bus network.

Because all of the star networks are connected to each other via the bus network, a break in the bus network disconnects the star networks from each other. However, a break in the bus does not affect communication within any of the individual star networks.

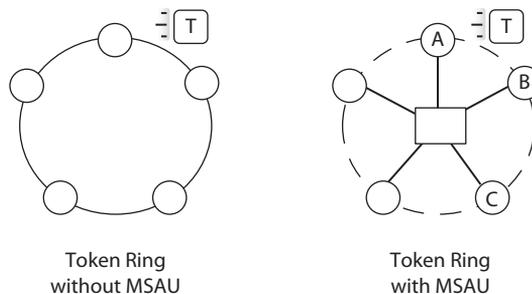
Token Ring

IEEE 802.5 defines token ring networks. In a token ring network, each device is connected together in a ring. Additionally, the ring includes a logical token that controls when computers can communicate. This prevents collisions because computers can send traffic on the network only when they have the token.

For example, consider a group of five people in a room sitting in a circle. Each person can talk, but only while holding a piece of paper identified as a token. Each person takes turns talking, and when they are done, they pass the token to the next person in the ring.

Figure 3-6 shows two diagrams of a token ring. The diagram on the left shows what the original token ring networks looked like. Each computer is connected to two other computers in the ring, and the computers pass the token around the network. If a computer has the token, it can send traffic on the network. However, if any of the computers fails

Figure 3-6
Token ring
configuration



or is disconnected from the network, it breaks the ring. No other computers will be able to communicate.

The configuration on the right includes a multistation access unit (MSAU), also called a media access unit (MAU). The traffic still flows logically in a circle, but it goes through the MSAU. For example, Computers A, B, and C aren't connected directly to each other, but instead they are connected to the MSAU. Traffic passes from Computer A to the MSAU, then to Computer B, back to the MSAU, and then to Computer C. If a computer fails, or a connection between the MSAU and a computer breaks, the MSAU senses the failure and bypasses the computer. In other words, a single failure does not take down the entire network.

Token ring networks don't scale well. Adding more computers significantly reduces performance. If a token ring has five computers, each computer has the token about 20 percent of the time. However, if the token ring has 100 computers, each computer has the token for only about 1 percent of the time. Traffic can slow to a crawl.

An extension of token ring networks is the Fiber Distributed Data Interface (FDDI) standard. It uses fiber-optic connections instead of copper connections and thus supports much higher bandwidth. It also uses dual rings. The second ring is used only when the system detects an error in the primary ring. Conventional token ring networks support speeds up to 16 Mbps. However, FDDI token rings support speeds up to 100 Mbps over distances as far away as 120 miles.

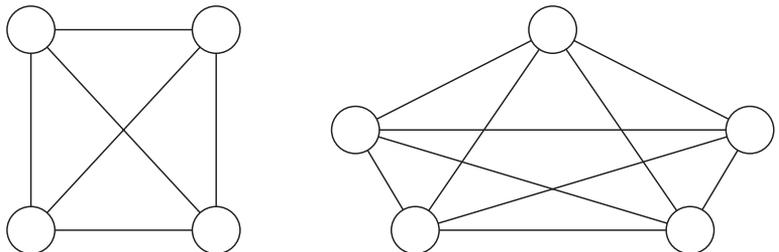
Mesh

Mesh topologies provide redundancy with multiple connections, providing the highest availability of any of the topology methods. A full mesh topology connects every computer or node in a network with every other node. A partial mesh topology provides multiple redundant connections but doesn't connect every node with every other node in the network. In either a full mesh or a partial mesh, computers in the network are still able to communicate with each other, even if some connections fail.

Figure 3-7 shows two examples of full mesh topologies. The one on the left has four nodes and six connections. The one on the right has five nodes and ten connections.

Notice how adding a single node requires the addition of multiple new connections. You can calculate the number of connections needed in a mesh topology with the formula

Figure 3-7
Mesh topology



$n(n-1)/2$, where n is the number of computers. For example, with four computers, $n=4$ and the formula is

- $4(4-1)/2$
- $4 \times 3/2$
- $12/2 = 6$

Add another computer and the calculation is $5(5-1)/2$, or 10 connections. Six computers require 15 connections. Each of these connections adds costs, so you'll rarely see full mesh topologies in a regular network. However, there are many instances where designers combine mesh topologies with another topology to create a partial mesh. This partial mesh topology has multiple connections to provide a high level of redundancy, but it doesn't connect every single computer with every other computer in the network.

Reviewing Basic Protocols and Ports

The TCP/IP protocol suite consists of a full range of protocols used to connect computers on the Internet and on most internal networks. You don't need to know the details of all the protocols, but you should understand the basic function of many of them. The following sections list some common protocols along with their acronyms and a brief explanation of the protocol.

It's also important to know the logical port numbers for many of the protocols. The "Mapping Well-Known Ports to Protocols" section later in this chapter describes the use of these ports in more detail and includes a table mapping some well-known ports to protocols. However, I've also listed the relevant port numbers in the protocol descriptions in case you turn to only that section for a description of the protocol.

A core step in hardening a server, or making it more secure from the default configuration, is to remove all unneeded protocols and services. With this in mind, it's important to know what protocols are required and what protocols are optional. Consider a server that is *not* running FTP. If an attacker tries to attack this server using a known FTP exploit, it simply won't work because FTP is not installed on the server.

On the other hand, consider a server that is running 100 different protocols. It is now susceptible to attacks on each of these protocols. By removing unneeded protocols, you reduce the attack surface of the server. Once you understand some basics about protocols, you can determine if a server needs specific protocols or if you can disable or remove them.

Figure 3-8 identifies the layers of the TCP/IP and OSI models where several protocols reside. The following pages include descriptions of these protocols.

Comparing IPv4 and IPv6

IP is the protocol used for addressing. Currently, both IPv4 and IPv6 addresses are in use on the Internet and in many internal networks. An IPv4 address uses 32 bits, but is typically displayed in dotted decimal format (such as 192.168.1.1), where each decimal number represents 8 bits.

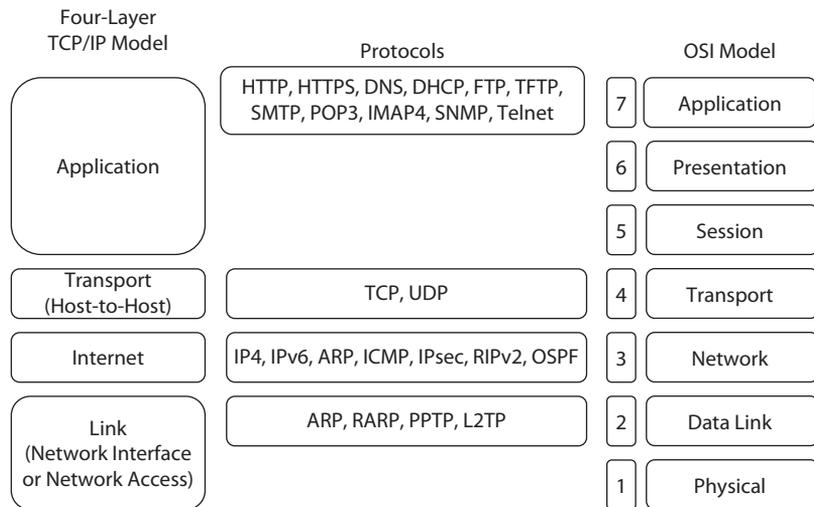


Figure 3-8 Mapping protocols to the OSI and TCP/IP models

An IPv6 address uses 128 bits, and is displayed with hexadecimal characters instead of decimal numbers. When the full IPv6 address is displayed, it includes eight groups of four hexadecimal characters separated by seven colons like this:

```
fe80:0000:0000:0000:045d:2376:095f:0075
```

Two methods used to shorten the displayed address include omitting leading zeros and using zero compression. When omitting leading zeros, the address looks like this:

```
fe80:0:0:0:45d:2376:95f:75
```

Even though this omits leading zeros, it's just for display. The actual IPv6 address still includes 128 bits.

Zero compression replaces a contiguous string of zeros with a double colon like this:

```
fe80::45d:2376:95f:75
```

Even though you don't see seven colons, or eight groups, this is also just for display. The actual IPv6 address still includes 128 bits.

Dynamic Host Configuration Protocol

A Dynamic Host Configuration Protocol (DHCP) server uses DHCP to assign TCP/IP configuration information to DHCP clients. This includes IP addresses, subnet masks, addresses of DNS servers, and much more. Many people shorten the definition of DHCP to say that it simply assigns IP addresses, but that is not completely accurate. There are close to 70 different options that a DHCP server can assign.

DHCP uses four primary messages between the client and the DHCP server:

- **DHCPDiscover** The DHCP client broadcasts this message when it is looking for a DHCP server that can provide the IP address and other configuration information.
- **DHCPOffer** DHCP servers that have IP addresses to offer for the client's subnet respond by broadcasting an offer. The message includes information such as the offered IP address, subnet mask, the IP address of the DHCP server making the offer, and any additional information configured on the DHCP server.
- **DHCPRequest** The DHCP client responds to the first offer it receives with a broadcast request to use the offer. This message includes the IP address of the DHCP server that made the offer.
- **DHCPAck** The DHCP server responds with a DHCP acknowledge message. At this point, the DHCP server re-serves that IP address to the DHCP client. When the DHCP client receives the information, it configures itself with the IP address and other TCP/IP configuration information.

DHCP uses UDP ports 67 and 68. The DHCP client sends traffic out on port 67 and receives answers on port 68. The DHCP server sends traffic on port 68 and receives traffic on port 67.



NOTE DHCP isn't limited to only running on a server. Many wireless routers include the DHCP service, which greatly simplifies the configuration of a wireless network.

Address Resolution Protocol

The Address Resolution Protocol (ARP) resolves IP addresses to MAC (or physical) addresses. The MAC address is a 12-hexadecimal number (such as 6C-62-6D-BA-73-6C) assigned to the NIC that uniquely identifies the computer. Systems use the MAC address when connecting to other computers on the same subnet.

From a big picture perspective, systems use the IP address to get packets from one subnet to another on the Transport layer. Once the packet arrives on the destination subnet, systems use the MAC address to get the frame to the destination computer on the Data Link layer.

ARP broadcasts a datagram that essentially asks all the computers on the subnet, "Who has this IP? If it's you, tell me your MAC address." The computer with that IP address responds with its MAC address.

Reverse Address Resolution Protocol (RARP) works just the opposite of ARP. A client has a MAC address and uses RARP to get an IP address. You can think of this as a simplified version of DHCP. It provides the client with some TCP/IP configuration information, but doesn't have the full capabilities of DHCP.

The Bootstrap Protocol (BootP) is similar to RARP in that it provides clients with IP addresses. However, it goes beyond just getting an IP address and can be used to retrieve a bootable image for diskless clients.



EXAM TIP ARP resolves IP addresses to physical or hardware addresses (MAC addresses). RARP allows a client with a MAC address to get an IP address. BootP allows a diskless client with a MAC address to get an IP address and retrieve a bootable operating system image.

Network Discovery Protocol

IPv6 uses the Network Discovery Protocol (NDP) instead of ARP to resolve IPv6 addresses to the IPv6 version of MAC addresses. IPv4 uses a 12-hexadecimal number (48 bits) for MAC addresses. IPv6 uses a 64-bit modified EUI-64 address instead of the IPv4 MAC address. EUI is an acronym for Extended Unique Identifier.

NDP does more than just resolve IPv6 addresses to EUI-64 addresses. Some of the tasks include

- Automatically configures nodes
- Discovers other nodes on the network
- Locates available routers on the network
- Detects duplicate addresses on the network

Domain Name System

A Domain Name System (DNS) server provides name resolution services. More specifically, DNS resolves host names to IP addresses as its primary function. For example, a DNS client can query the DNS server with the host name of a computer, and the DNS server replies with the host's IP address.

DNS uses a hierarchical naming system with a distributed database to provide responses to clients. Any single DNS server doesn't know the IP address of all the systems on the Internet. However, by querying different servers in the hierarchy, it's possible to discover the IP address of any system with a public IP address.

Thirteen root servers at the top of the hierarchy know the IP addresses for top-level domain DNS servers. Top-level domains include .com, .gov, .edu, .net, and so on. These top-level domain DNS servers know the IP address of second-level DNS servers within the same domain. For example, a .com DNS server knows the IP address of authoritative DNS servers in the mcgraw-hill.com domain. Second-level domain DNS servers know the IP address of third-level DNS servers within the same domain.

The distributed database includes several different types of records to provide different responses. Some of the records in a DNS database are

- **A (host) record** Resolves a host name to an IP address
- **PTR (pointer) record** Resolves an IP address to a host name

- **MX (mail exchange) record** Identifies a mail server
- **CNAME (canonical name or alias) record** Allows a system to be known by different names, or aliases. For example, a web server can be named Web1 but also respond to a name of www.

Berkeley Internet Name Domain (BIND) is a version of DNS software that runs on UNIX systems. It is freely available and runs on many DNS servers on the Internet. Microsoft networks running in internal organizations commonly use Microsoft's version of DNS.

When clients query DNS for name resolution, DNS uses UDP port 53. However, when DNS servers transfer data between each other, they use TCP port 53.



EXAM TIP DNS uses a hierarchical structure with a distributed database to provide name resolution services. BIND is commonly used on the Internet, while Microsoft DNS is often used within a Microsoft internal network.

Internet Control Message Protocol

The Internet Control Message Protocol (ICMP) is used with many diagnostic protocols such as ping, pathping, and tracert to check or verify the health of a network. For example, if you enter **ping** and an IP address (such as **ping 192.168.1.10**) on a Windows system, ping sends four packets to the IP address. If firewalls aren't blocking ICMP, the host with that IP address will reply with four ping packets. Other operating systems, such as Linux and UNIX, will continue to send ping requests until the user presses CTRL-C.



TIP UNIX and Linux systems use traceroute instead of tracert. While there are many similarities between the two commands, traceroute uses UDP instead of ICMP.

Attackers often use ICMP in DoS attacks, so firewalls frequently block ICMP traffic. In other words, a server may be operational, but a ping to that server will fail due to a firewall blocking ICMP traffic.

Internet Group Message Protocol

The Internet Group Message Protocol (IGMP) is used for IPv4 multicasting, where data is sent from one computer to multiple computers at the same time. Computers join a multicast group with a multicast address (in the range of 224.0.0.0 through 239.255.255.255) to receive the multicast transmissions.



TIP IPv6 does multicasting differently than IPv4. IPv6 uses Multicast Listener Discovery (MLD) with ICMPv6 messaging instead of IGMP.

You can compare multicasting used with IGMP to unicast and broadcast messages:

- **Unicast** Systems send traffic from one computer to one other computer.
- **Broadcast** Systems send traffic from one computer to all other computers on the subnet. Switches pass broadcast traffic, but routers do not pass broadcast traffic, so broadcast traffic only goes to computers on the same subnet.
- **Multicast** Systems send traffic from one computer to multiple computers in the network. Routers pass multicast traffic, so multicast transmissions can reach computers on multiple subnets within the network. Many audio and video streaming applications use multicast, and multicast uses UDP.

Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) is commonly used to manage network devices (such as routers and layer 3 switches) on a network. SNMP agents on the network devices send traps (errors) and other notifications back to a central server running SNMP management software. Additionally, a server hosting SNMP management software can use SNMP to configure the network devices remotely.

SNMPv1 and SNMPv2 had several security flaws, but SNMPv3 provides significant improvements. For example, client authentication in SNMPv1 was handled with a community string, and the default community was “public” in SNMPv1 and SNMPv2. However, SNMPv3 introduced several security enhancements. SNMPv3 uses encryption to achieve confidentiality, verifies the integrity of the packets, and includes authentication to verify that the data is from a known entity. SNMP agents receive data on UDP 161 and send traps and notifications on UDP port 162. When using Transport Layer Security (TLS) to encrypt data, Secure SNMP uses TCP ports 10161 and 10162.

When SNMP is active on a network, internal routers will pass SNMP traffic. However, it’s common to block all SNMP traffic at the boundary of the private network with the Internet. This prevents attackers from accessing internal devices using SNMP.

File Transfer Protocol

File Transfer Protocol (FTP) is used to transfer files from one system to another. Some organizations host FTP servers that allow individuals to upload and download data. The FTP server provides a simple storage location for people located in different areas of the world. One benefit of FTP is that it supports authentication in the form of a username or password, but administrators can also configure FTP to support anonymous connections. FTP uses TCP ports 20 and 21.

Trivial FTP (TFTP) was mentioned briefly earlier in the chapter in the “UDP” section. TFTP is a stripped-down version of FTP that provides very basic capabilities. It doesn’t support authentication and doesn’t use TCP to establish a session. However, it is very useful when transferring configuration files to and from network devices. TFTP uses UDP port 69.



EXAM TIP FTP supports authentication and is commonly used to transfer large files across a network. FTP uses TCP ports 20 and 21. TFTP does not support authentication and is commonly used to transfer small configuration files to and from network devices. TFTP uses UDP port 69. Both transfer data across a network in cleartext.

Both FTP and TFTP have significant security issues. Data sent across the network using FTP and TFTP is sent in cleartext. Anyone with a sniffer can easily capture the data and read the contents of the FTP packets. This includes the username and password used to log on to an FTP server, and the data.

Secure FTP (SFTP) uses Secure Shell (SSH) to encrypt FTP traffic. SFTP functions similarly to FTP, but the encryption prevents sniffing attacks from capturing data in cleartext.

Telnet

Telnet is a basic protocol used to interact with a remote system using text-based commands. For example, a Telnet server will accept connections from a system running a Telnet client. When the client connects, the user can enter commands from a Telnet command prompt. Telnet uses TCP port 23.

Although Telnet used to be much more popular, especially in UNIX and UNIX derivatives such as Linux, you won't see it in use as often today. Telnet sends data (including login credentials) across the network in cleartext, which attackers can capture and read using a sniffer. Once attackers have the credentials, they can log on to the Telnet server and launch additional attacks.

Secure Shell

Many Linux and UNIX systems support rlogin (remote login), a utility that allows users to log in to a remote computer. Similarly, many Linux and UNIX systems support rsh (remote shell) and rexec (remote execute), which allow users to execute commands on remote computers. However, these commands send data across a network in cleartext, representing significant risks.

Secure Shell (SSH) is an encryption protocol that can create a secure session between two computers. It is used to encrypt other protocol traffic, such as FTP (called Secure FTP or SFTP) and Secure Copy (SCP), and can replace many protocols such as Telnet, rlogin, rsh, and rexec. SSH also supports secure port forwarding. SSH uses TCP port 22.

SSH provides several security services. It uses asymmetric and symmetric encryption to provide confidentiality, and hashing techniques to provide integrity. SSH also supports mutual authentication, where both the client and the server prove their identity within an SSH authenticated tunnel. RFCs 4251, 4252, 4253, and 4254 document the SSH protocol.



EXAM TIP SSH can encrypt other protocols such as FTP (as SFTP). It is a more secure alternative than older protocols such as Telnet, rlogin, rsh, and rexec that send data across a network in cleartext. SSH uses a combination of asymmetric and symmetric encryption for confidentiality, hashing for integrity, and an SSH authenticated tunnel for authentication. SSH uses TCP port 22.

HyperText Transfer Protocol and HyperText Transfer Protocol Secure

HyperText Transfer Protocol (HTTP) is the common protocol used to transfer web pages over the Internet. Most web browsers use HTTP by default. For example, you could type in the address of a site such as www.mcgraw-hill.com and it will default to <http://www.mcgraw-hill.com>. HTTP uses TCP port 80.



NOTE Some web browsers hide “http://” in the address bar so you may not see it. However, if you copied the URL and pasted it into a document, you would see that it is including <http://> even if you don’t type it in the address bar.

HTTP Secure (HTTPS) is an encrypted form of HTTP and is commonly used when processing e-commerce transactions over the Internet. Secure Sockets Layer (SSL) encrypted HTTPS originally, but Transport Layer Security (TLS) is the designated replacement for SSL and has replaced it in almost all applications. Data encrypted within HTTPS is not readable by anyone who is not involved in the HTTPS session. HTTPS uses TCP port 443.

Transport Layer Security and Secure Sockets Layer

Transport Layer Security (TLS) is a popular encryption protocol used to encrypt different types of traffic. You most often see TLS used to encrypt HTTP traffic (as HTTPS). However, other implementations also use TLS. For example, some VPNs use TLS to encrypt tunnel traffic.

Transport Layer Security (TLS) is the designated replacement for Secure Sockets Layer (SSL), and is based on SSL. Although the acronym HTTPS historically has meant that SSL is encrypting an HTTP session, many current HTTPS sessions use TLS. Chapter 14 describes the process used by TLS and SSL when encrypting HTTPS traffic.



TIP SSL is vulnerable to POODLE (Padding Oracle On Downgraded Legacy Encryption) attacks and no longer recommended for use. Personnel from Google documented this attack in late 2014, which prompted organizations still using SSL to switch to TLS.

IT personnel often debate the actual OSI layer where these protocols operate. RFCs typically provide a definitive answer, but both SSL and TLS are layered protocols, meaning that they operate on different layers. RFC 6101 is a historic RFC that describes SSL.

RFCs 2246, 4346, 5246, and 6176 formally define TLS. The following list describes operations performed by TLS on different layers:

- **Transport layer** TLS creates the initial connection on the Transport layer (layer 4). With the name Transport Layer Security, it's implied that TLS operates on the Transport layer.
- **Session layer** The Handshake Protocol within TLS operates on the Session layer (layer 5) when negotiating details for the TLS session. This includes details such as what specific encryption algorithm to use.
- **Application layer** The TLS Record Protocol within TLS operates on the Application layer (layer 7). It provides confidentiality with an encryption algorithm and integrity with a hashing algorithm.

Network File System

Network File System (NFS) allows computers running different operating systems to access and share files over the network. NFS is an open standard created by Sun Microsystems. It allows users running UNIX (and UNIX derivatives such as Linux) to access files on Microsoft systems. Similarly, NFS allows users running Microsoft systems to access files on UNIX-based systems.

Routing Protocols

As mentioned previously, routing occurs at the Network layer. Routers identify the best path for a packet to take based on information it knows about the network. It's possible to enter routing information manually in a smaller network, but for larger networks, the routers communicate with each other using a routing protocol. The two primary routing protocols used on internal networks are Routing Information Protocol version 2 (RIPv2) and Open Shortest Path First (OSPF).

Although RIPv2 works well, it can generate a lot of traffic and doesn't scale well. In general, RIPv2 is good for smaller networks, while OSPF works better in larger networks. RIPv2 is very easy to enable and configure on routers, while OSPF takes a little more effort. Both RIPv2 and OSPF operate on the Network layer of the OSI Model.

E-mail Protocols

You should know about three primary e-mail protocols when preparing for the SSCP exam: Simple Mail Transfer Protocol (SMTP), Post Office Protocol version 3 (POP3), and Internet Message Access Protocol version 4 (IMAP4).

Clients use SMTP to send e-mail to an e-mail server for delivery. Some clients receive e-mail from a POP3 server, and other clients access e-mail from an IMAP4 server. POP3 servers simply send the e-mail to the client when the client connects. Some POP3 servers delete the e-mail as soon as the client receives it. Other POP3 servers keep the e-mail for a preset time limit, or until the user deletes it off the server. IMAP4 servers store e-mail on the server and allow users to organize the e-mail in folders. Another advantage of IMAP4 over POP3 is that users can configure it to show only the headers. This allows

users to screen some of the e-mail before downloading it to their system. In contrast, POP3 sends the entire contents to the client as soon as the client connects.

SMTP uses TCP port 25, POP3 uses TCP port 110, and IMAP4 uses TCP port 143. Each of these protocols supports encryption with SSL or TLS. SMTP with SSL/TLS uses TCP port 465, IMAP4 with SSL/TLS uses TCP port 993, and POP3 with SSL/TLS uses TCP port 995.

With the overwhelming amount of spam and phishing attacks, IMAP4 is an appealing alternative used by many e-mail servers on the Internet. Users can often tell by the sender e-mail address or the subject line whether they want to open the e-mail or just delete it. For example, I'd love to actually win \$1,500,000 just because I have a Gmail account, but when I see the subject line saying, "Your e-mail ID has won \$1,500,000," I know it's a phishing attempt and I simply delete the message to eliminate the risk.



EXAM TIP Users send e-mail with SMTP over TCP port 25. Users receive e-mail with POP3 over TCP port 110 or IMAP4 using TCP port 143. IMAP4 allows users to manipulate the e-mail on the server, including organizing e-mail in folders.

Tunneling Protocols

Chapter 4 explores virtual private networks (VPNs) in greater depth. In general, a VPN provides access to a private network over a public network such as the Internet. Because traffic traverses a public network, potential attackers can access the data using a protocol analyzer or sniffer. A common method of protecting VPN traffic is to encapsulate it in a tunneling protocol and encrypt the encapsulated data.

Two common tunneling protocols are the Point-to-Point Tunneling Protocol (PPTP) and the Layer 2 Tunneling Protocol (L2TP). PPTP provides encryption, while L2TP relies on IPsec for encryption. Both operate on the Data Link layer. PPTP uses TCP port 1723 and L2TP uses UDP port 1701. Chapter 4 explores these protocols in more depth.

Internet Protocol Security

Internet Protocol security (IPsec) provides security for IP traffic traveling over a network. It provides mutual authentication of systems, integrity, and confidentiality. Chapter 4 expands on IPsec, but as an introduction, IPsec has two primary components:

- **Authentication Header (AH)** AH provides authentication between the systems and verifies the integrity of the packets, but it does not encrypt the data. The header content is created from a hash of the packet with additional authentication data, and this hash is then encrypted to prevent tampering. The hash provides integrity. Systems identify AH packets with protocol number 51.
- **Encapsulating Security Protocol (ESP)** ESP encrypts the data and includes the same authentication services provided from AH. Data encrypted with ESP cannot be read if captured with a sniffer. Systems identify ESP packets with protocol number 50.



EXAM TIP Note that AH and ESP are identified with protocol numbers (not port numbers). The “Comparing Ports and Protocol Numbers” section a bit later in this chapter identifies the differences. The IP header includes the protocol number field. Other protocol number values are as follows: ICMP is 1, IGMP is 2, TCP is 6, and UDP is 17.

Mapping Well-Known Ports to Protocols

There are a total of 65,536 TCP ports and 65,536 UDP ports. They are divided into three sections:

- **Well-known ports (0 to 1023)** These ports are assigned to widely used protocols by the Internet Assigned Numbers Authority (IANA).
- **Registered ports (1024 to 49,151)** These ports are also assigned by IANA, although most are not as widely used as the well-known ports.
- **Dynamic (49,152 to 65,535)** These ports are assigned dynamically by internal systems.

As an example of how ports are used, consider Figure 3-9. The client is using a web browser to connect to an organization’s internal web server. When the user clicks a link to connect to the web server, several things happen to retrieve a web page. To begin, the client creates a datagram and sends it to the web server.

The datagram includes the source and destination IP addresses and the source and destination ports. A web server has an IP address of 192.168.1.5 and uses TCP port 80 for HTTP, so the destination port is port 80. The client adds its IP address (192.168.1.20) as the source IP address and then identifies an unused port in the dynamic range and assigns it as the source port for this datagram. The user’s computer maps the source port to the user’s web browser application.

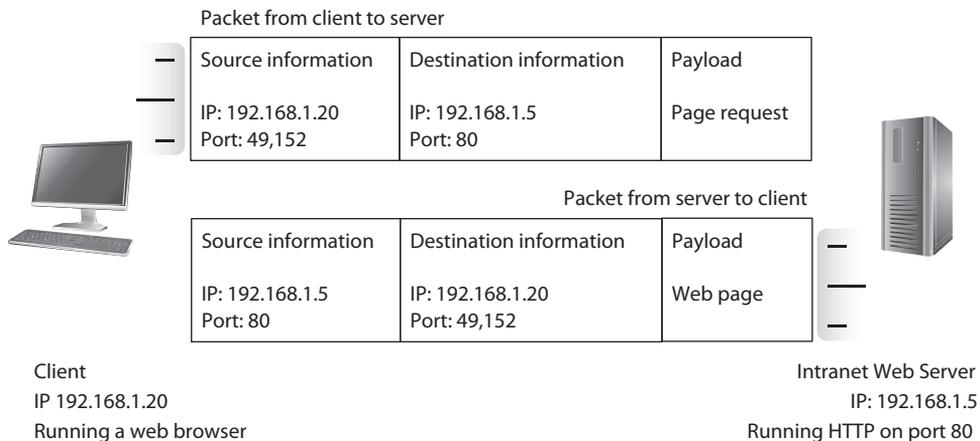


Figure 3-9 Using IP addresses and ports

TCP/IP gets the datagram to the server using the destination IP address. Once the datagram reaches the server, the server uses the destination port to get the datagram to the correct protocol, service, or application that is servicing that port. In this example, the destination is port 80, which is the well-known port for HTTP. The server sends the datagram to the web server application servicing HTTP.

The web server then creates the web page and puts the data together to return the web page to the client. Because the source of the original packet was 192.168.1.20 with a port of 49,152, this now becomes the destination. Again, TCP/IP gets the packet to the client computer using the destination IP address.

When the client receives the packet, it identifies the destination port (49,152), which it previously mapped to the web browser application. The operating system sends the data to the web browser, which then displays it.

Table 3-1 lists many common protocols and ports, along with some basic comments about the protocol.

Protocol	Port(s)	Comments
FTP	TCP 20, 21	Optional but used on FTP servers. Data is sent in cleartext but can be secured with SSH.
TFTP	UDP 69	Optional and usually disabled.
SSH	TCP 22	Optional. It is used to secure other protocols, such as FTP (SFTP) and Secure Copy (SCP).
Telnet	TCP 23	Optional and usually disabled. Data is sent in cleartext, and Telnet can be used to launch other attacks. SSH is a preferred alternative.
DNS	TCP 53 UDP 53	Required for DNS. Users send DNS queries using UDP 53. DNS servers transfer data using TCP 53.
DHCP	UDP 67, 68	Needed on DHCP servers and for DHCP clients. Clients send DHCP data on UDP 67 and receive DHCP data on UDP 68.
SMTP	TCP 25	Used to send e-mail. E-mail servers and clients with e-mail applications need SMTP enabled. Some web servers also use SMTP.
POP3	TCP 110	Used to receive e-mail from a POP3 server.
IMAP4	TCP 143	Used to store, manipulate, and manage e-mail on an e-mail server.
HTTP	TCP 80	Needed on web servers.
HTTPS	TCP 443	Needed on web servers that support secure transmissions. Uses TLS or SSL.
SNMP	UDP 161, 162	Used to manage network devices (such as routers and layer 3 switches). Agents receive data on UDP 161 and send traps and notifications on UDP 162.
Secure SNMP	TCP 10161, 10162	Secured with Transport Layer Security.
Kerberos	TCP 88	A network authentication protocol (presented in Chapter 2).
PPTP	TCP 1723	Used with VPNs.
L2TP	UDP 1701	Used with VPNs.

Table 3-1 Common Protocols and Their Port Numbers

Table 3-2
Common Protocol
Numbers

Protocol Number	Protocol
1	ICMP (IPv4)
2	IGMP (IPv4)
6	TCP
17	UDP
50	IPsec ESP
51	IPsec AH
56	ICMP (IPv6)

Comparing Ports and Protocol Numbers

While some documentation indicates that ports are the same as protocol numbers, they are completely different. The first octet in TCP and UDP headers includes source and destination ports for TCP and UDP segments, respectively. As a reminder, TCP and UDP operate on the Transport layer of the OSI Model. IPv4 packets include protocol numbers in the Protocol field of the IPv4 header, and IPv6 packets include protocol numbers in the Next Header field of the IPv6 header. As a reminder, IPv4 and IPv6 operate on the Network layer of the OSI Model.

Table 3-2 lists common protocol numbers.



EXAM TIP Know the common ports and protocol numbers listed in Tables 3-1 and 3-2. Routers and firewalls use these to control traffic in and out of a network, so these are important security concepts to know.

Comparing Internetwork Trust Architectures

Computers are connected together using networks, and different types of networks provide different levels of trust. Generically, there are four types of networks: the Internet, an intranet, an extranet, and a demilitarized zone (DMZ, or perimeter network). Each of these can support different types of trust architectures, such as one-way trusts, two-way trusts, and transitive trusts.

Consider Figure 3-10 as you review the definitions of the first three different network types in the following list:

- **Internet** The Internet is the public network accessible around the world. Any computer with access to the Internet can reach any other computer with access to the Internet. From a risk perspective, the Internet presents the most risk.
- **Intranet** An organization's internal network is an intranet. From a risk perspective, an intranet presents the least amount of risk. An organization controls the resources on the internal network.

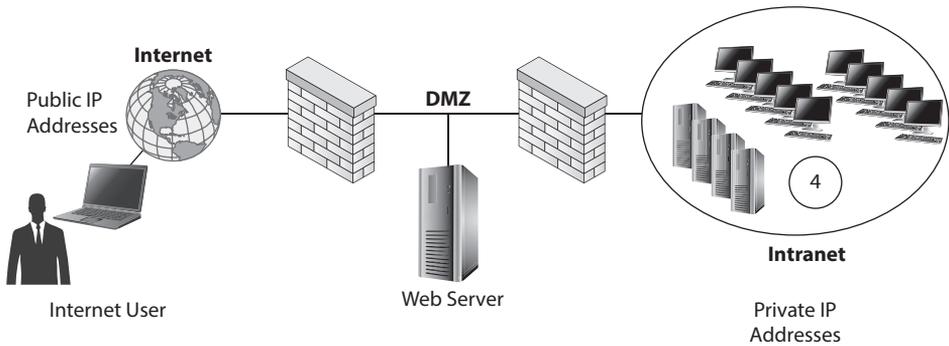


Figure 3-10 Comparing the Internet, an intranet, and a DMZ

- DMZ** When an organization wants to host resources on the Internet (such as with a web server, an e-mail server, or an FTP server), it often places these Internet-facing servers within a DMZ. The DMZ provides a layer of protection for the servers that would not be available if they were placed directly on the Internet. Although there are multiple methods for creating a DMZ, a common method is to use two firewalls, as shown in Figure 3-10. Another method is to create a multi-homed firewall, using three or more NICs to create a separate screened subnet.
- Extranet and third-party connections** An extranet is a network used to host resources via the Internet, but is accessible only to trusted third-party entities, as shown in Figure 3-11. For example, a company may have a partnership with another company and want to share resources with that partner through a web-based portal. Administrators at the first company can create an extranet to provide access to these resources, but restrict access to this network to only the trusted partner. Many organizations use federated access systems to provide single sign-on (SSO) capabilities to systems hosted in an extranet. Chapter 2 describes federated access in more depth.

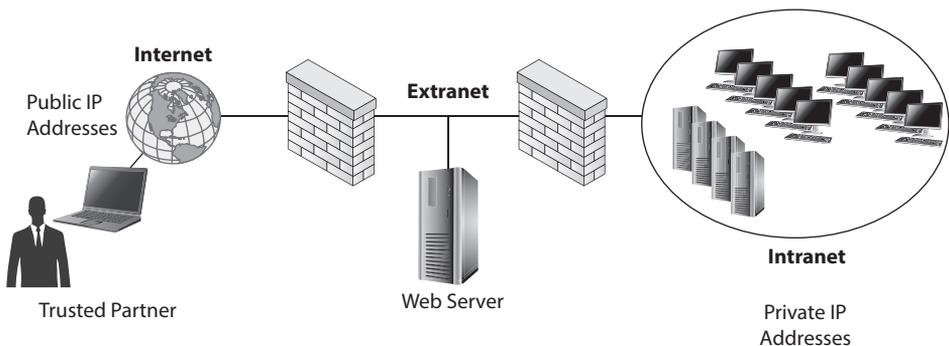


Figure 3-11 Comparing the Internet, an intranet, and an extranet



TIP Notice that the drawing representing an extranet (Figure 3-11) is very similar to the drawing for a DMZ (Figure 3-10). The difference is in the intent. Resources in a DMZ are available to any users on the Internet. Administrators configure an extranet so that resources are only available to trusted entities. An organization can host both a DMZ and an extranet.

Additionally, networks are identified with some other terms as follows:

- **Personal area network (PAN)** A small network centered on a person. For example, a person may have a cell phone paired with an earpiece and a microphone within a PAN.
- **Local area network (LAN)** A network that connects computers together in the same geographical location. A LAN is a single, well-connected network. It could be in a single room, a single building, or a group of buildings connected together with links of similar speeds.
- **Metropolitan area network (MAN)** A group of networks that are connected over a large area such as a university campus or even an entire city.
- **Wide area network (WAN)** A large network that connects computers across a large geographical area, such as in different cities. WAN links connect WANs together, and WAN links are typically slower than the links within individual LANs.

Comparing Public and Private IP Addresses

You may have noticed in Figures 3-10 and 3-11 that the Internet included public IP addresses, while the intranet included private IP addresses. Public IP addresses are used only on the Internet, and private IP addresses are used within an organization. A public IP address is accessible from any other public IP address. In other words, any computer (or any attacker) with access to the Internet from anywhere in the world can access another computer on the Internet because they both have public IP addresses.

Every public IP address must be unique. No two computers can be assigned the same public IP address. In contrast, private IP addresses must be unique only within an organization. For example, all the IP addresses used by McGraw-Hill in its network are unique within the McGraw-Hill network. However, (ISC)² can use the same private IP addresses as McGraw-Hill because both companies use these IP addresses internally.

It's common for IP addresses within the DMZ or within an extranet to have public IP addresses. However, it's also possible to configure the firewalls to forward traffic to servers within a DMZ or extranet and configure these servers with private IP addresses.

RFC 1918 lists the IPv4 addresses reserved for private networks. They are as follows:

- **10.0.0.0 through 10.255.255.255** A very large address space with enough addresses for large organizations
- **172.16.0.0 through 172.31.255.255** A medium-sized address space
- **192.168.0.0 through 192.168.255.255** A smaller address space for smaller organizations

Link Local Addresses

IPv4 and IPv6 both have specific IP address ranges that can be assigned to clients automatically. These are called link local addresses. A computer with only a link local address can communicate with other computers on the local link or local subnet, but not with computers on any other subnets. In other words, these addresses cannot pass a router.

The link local address range in IPv4 is 169.254.0.1 through 169.254.255.254, more commonly known as Automatic Private IP Addresses (APIPAs). DHCP clients will automatically assign themselves an address in this range if they do not receive an IP address from a DHCP server.

IPv6 works a little differently. IPv6 clients will typically assign themselves a link local address any time IPv6 is enabled. It doesn't matter if they've received an IPv6 address from a DHCP server, or even if they have a manually assigned IPv6 address. The IPv6 link local address range starts with the `fe80::/10` prefix.



EXAM TIP Private IP addresses are not routable on the Internet. If traffic does reach the Internet with a private source or destination IP address, routers on the Internet drop it.

RFC 4193 lists the IPv6 addresses reserved for private networks. They are identified as unique local IPv6 unicast addresses (or local unicast addresses). They have a prefix of `fc00::/7`, which is a little misleading.

The `/7` in `fc00::/7` identifies the first seven bits in the address. A hexadecimal `f` is binary `1111`, and a hexadecimal `c` is binary `1100`. If only the first seven characters are used, it is `1111 110x`, with `x` being the eighth bit. However, RFC 4193 specifies that the eighth bit is always a `1` (although this may change in the future), so the first eight characters of a unique local address are `1111 1101` or `fd`. In other words, it's technically accurate to say that the link local address starts with `fc00::/7`, but in practice, you'll actually see link local addresses that start with `fd` (using a hexadecimal `d` instead of a hexadecimal `c`).

Using NAT

The Internet has public IPv4 addresses, and internal networks have private IPv4 addresses. However, users with private IPv4 addresses can still access the Internet. Network Address Translation (NAT) provides a mechanism that allows this to happen. NAT is an additional service running on a system at the boundary between the Internet and an internal network.

Consider Figure 3-12. It shows an intranet with private IP addresses as well as the Internet with public IP addresses, separated by a proxy server with NAT installed. Internal clients connect to the Internet via the proxy server, and NAT translates the IP addresses from public to private and from private back to public.

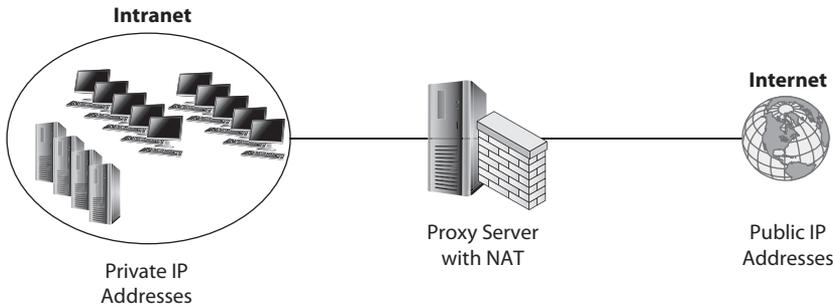


Figure 3-12 Using a proxy server with NAT

One of the benefits of NAT is that internal clients do not need public IP addresses to access the Internet. This provides a layer of protection for these clients because they aren't directly accessible from other Internet clients. It also reduces costs for an organization because the organization needs to purchase public IP addresses only for the proxy server and not for every internal client.

Smaller networks often use NAT on a router instead of a proxy server. For example, wireless routers used in small wireless networks often include NAT. All internal clients use private IP addresses and access the Internet via the wireless router.

NAT is typically identified as either static or dynamic:

- **Static** Maps an unregistered private IP address to a registered public IP address on a one-to-one basis. Static NAT allows an internal client to communicate with devices on the Internet with the same IP address.
- **Dynamic** Uses multiple public IP addresses and maps private IP addresses to one of the public IP addresses based on a first-come, first-served basis.

Port Address Translation (PAT) is an extension of NAT that maps multiple private IP addresses to a single public IP address using different ports. As an example, imagine the proxy server in Figure 3-12 is using PAT instead of NAT and Sally uses her web browser to connect to mcgraw-hill.com from the private network. The PAT server changes the request using the PAT server's public IP address as the source IP address, and it changes the source port to a unique source port number, such as 49,152. When the PAT server receives the response, it has a destination port of 49,152. The proxy server recognizes this port and modifies the packet using Sally's original private IP address and port number, and forwards it to her computer.



EXAM TIP Static NAT uses a one-to-one mapping with a single IP address, but dynamic NAT uses a one-to-many mapping with multiple IP addresses. PAT is an extension of NAT that uses a many-to-many mapping using port numbers.

Comparing Trust Relationships

Networks utilize trust relationships to simplify access control. The three primary types of trust relationships are one-way, two-way, and transitive. Refer to Figure 3-13 when reviewing the following explanations:

- One-way trust** A one-way trust allows subjects (such as users) in one domain to access resources in another domain. For example, Figure 3-13 shows a one-way trust relationship where mheducation.com trusts Training.McGrawHill.com. This would allow users in the Training.McGrawHill.com domain (such as Sally) to access resources in the mheducation.com domain. However, users in the mheducation.com domain would not be able to access resources in the Training.McGrawHill.com domain.



NOTE A one-way trust does not automatically grant permissions. In other words, everyone in the Training.McGrawHill.com domain would not automatically have access to all resources in the mheducation.com domain. However, the one-way trust makes it possible to grant users permissions to the resources.

- Two-way trust** A two-way trust indicates that both domains trust each other. In Figure 3-13, there is a two-way trust between the parent domain (McGrawHill.com) and each child domain (Training and Publishing). You can also think of this as two one-way trusts.
- Transitive trust** A transitive trust is an indirect trust relationship. In Figure 3-13, the Training domain and the Publishing domain do not have a direct trust relationship. However, because they both have a two-way trust relationship with

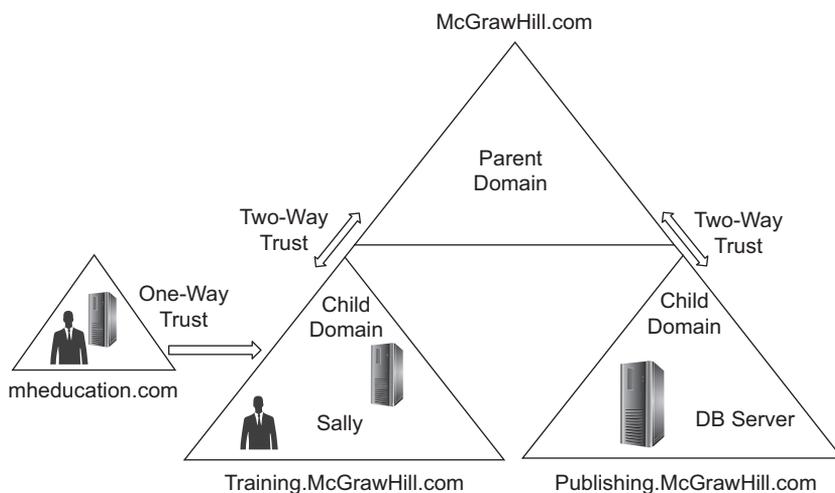


Figure 3-13 Trust relationships

the parent domain (McGrawHill.com), it creates a transitive trust relationship between both domains. This allows users in the Training domain to access resources in the Publishing domain, and it allows users in the Publishing domain to access resources in the Training domain. This doesn't mean they automatically have full rights and permissions in the other domain. However, administrators can grant users rights and permissions when the transitive trust relationship exists.

Exploring Wireless Technologies

Wireless technologies have become quite popular in the past decade. With a wireless network, computers connect to each other and network devices by transmitting data over the airwaves. There are five primary IEEE 802.11 standards currently in use. Table 3-3 outlines the base frequency and theoretical maximum speeds of each.



NOTE The speed is the maximum speed in ideal conditions. However, many variables affect the actual speed. For example, as the distance increases between two wireless devices, the actual speed decreases. Obstructions such as walls and interference from EMI or RFI also reduce the actual speed.

A wireless network includes a wireless access point (WAP) or a wireless router. Wireless devices, such as laptops, tablets, and smartphones, connect to the network through the WAP or the wireless router. A WAP provides a bridge for wireless clients to a wired network. For example, Figure 3-14 shows a WAP for an organization with one user connecting to the WAP with a smartphone while another user is connecting with a laptop. Once connected, the wireless clients are able to access any resources in the network, just as other wired clients can.

Wireless routers include additional capabilities beyond the WAP. Technically, they are access points with routing capabilities, but they are widely marketed as wireless routers. They are widely used in small offices and home offices (SOHOs), along with many home networks. Some of the common capabilities found in a wireless router include

- **WAP** Provides a bridge for wireless clients to the wired network.
- **DHCP** Provides both wired and wireless clients with TCP/IP configuration information, such as IP addresses, subnet masks, the address of a DNS server, and more.

Standard	Base Frequency	Maximum Speed
802.11a	5 GHz	54 Mbps
802.11b	2.4 GHz	11 Mbps
802.11g	2.4 GHz	54 Mbps
802.11n	2.4 GHz, 5 GHz	600 Mbps
802.11ac	2.4 GHz, 5 GHz	600 Mbps (2.4 GHz), 2.6 Gbps (5 GHz)

Table 3-3 Wireless Standards

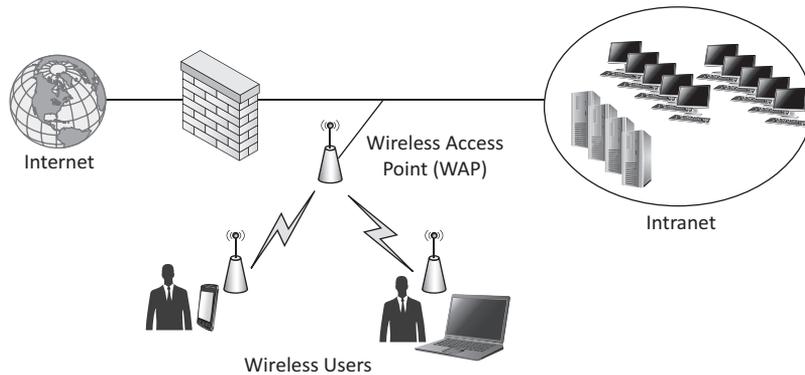


Figure 3-14 Wireless users connecting to a network via a WAP

- **NAT** Translates between private IP addresses used internally and public IP addresses used on the Internet. Wireless routers typically use PAT, which is an extension of NAT.
- **Routing** Provides basic routing between the internal network and the Internet. The wireless router is the default gateway for all internal clients.
- **Basic firewall** Provides basic firewall capabilities to filter traffic, but the capabilities can differ substantially from one brand to another.

Figure 3-15 shows an example of a wireless router within a network. This wireless router is providing many additional services for both wired and wireless clients. All clients access the Internet via this wireless router.

Wireless systems typically transmit in multiple directions (omnidirectional) at the same time. The area where wireless systems can receive the signal is the wireless footprint.

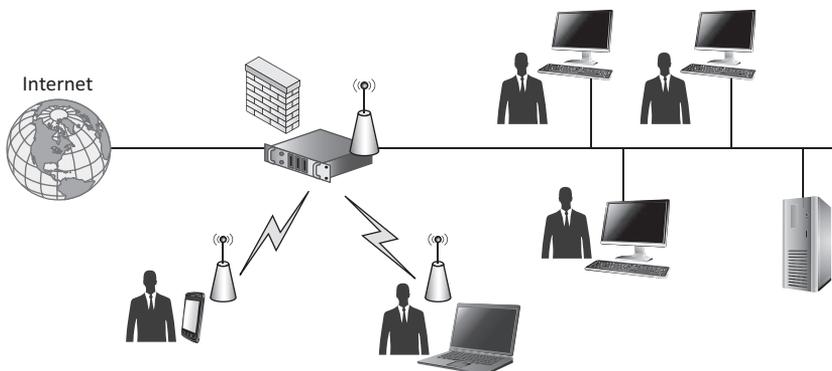


Figure 3-15 Wireless users connecting to a network via a wireless router

Within an organization, walls or devices transmitting EMI or RFI reduce the overall footprint.

However, an attacker can extend the footprint of a wireless network by using a receiver with a directional antenna. For example, it's possible to hook up something as simple as a can to a wireless antenna (creating a cantenna) and then point the can toward the wireless transmitter. Some tests have shown that an attacker can capture transmissions of wireless networks more than a mile away by using a cantenna.

Securing Data Transmissions

Because wireless technologies transmit data over the air using radio waves, it's easy for an attacker to capture the data with a wireless receiver. This can be as simple as installing a wireless sniffer on a laptop that includes wireless capabilities. For example, Wireshark is a free sniffer application that includes the ability to capture wireless transmissions. An attacker can download and install it, and then sit in an organization's parking lot and capture wireless transmissions.



TIP Wireshark can capture all transmissions as long as the NIC is in promiscuous mode. If it's not in promiscuous mode, Wireshark only captures traffic addressed to or from the Wireshark client. The challenge with capturing wireless transmissions is that many wireless NICs do not support promiscuous mode. However, some adapters such as AirPcap support promiscuous mode.

Secure wireless transmissions make it more difficult for an attacker to capture any usable information from these transmissions. The three primary security protocols are Wired Equivalent Privacy (WEP), Wi-Fi Protected Access (WPA), and Wi-Fi Protected Access 2 (WPA2).

WEP was the first security algorithm used for 802.11 wireless networks. The goal was to provide the same level of privacy on a wireless network as users could achieve on a wired network. It failed. WEP has several security issues and should not be used.

WPA was the interim replacement for WEP. The goal was to provide better security immediately for wireless devices using existing hardware. In other words, users running WEP could switch over to WPA without purchasing new wireless adapters, wireless access points, and/or wireless routers. Although WPA is much better than WEP, its designers never intended it to be a permanent replacement. WPA has known vulnerabilities and is susceptible to WPA cracking attacks allowing an attacker to discover the WPA preshared key (PSK), which is a password or passphrase of 8 to 63 characters.

WPA2 is the permanent replacement for WEP and WPA. The IEEE formally ratified it in 2004 as 802.11i. WPA2 requires different hardware than WEP and WPA. However, the Wi-Fi Alliance requires that all new hardware support WPA2 and Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (CCMP) in order to be Wi-Fi Certified. The Wi-Fi Alliance is a global nonprofit organization that has helped standardize wireless technologies.

WPA originally used only Temporal Key Integrity Protocol (TKIP). TKIP could use the same hardware as WEP, but provided significantly more security than WEP. Later implementations of WPA support Advanced Encryption Standard (AES), a much stronger encryption standard than TKIP.

WPA2 uses CCMP, which is based on AES and is much stronger than WPA using either TKIP or AES. WPA2 with CCMP is the recommended standard for security.

WPA-Personal and WPA2-Personal

WPA-Personal and WPA2-Personal use a preshared key (PSK). WPA/WPA2 then uses this PSK to create an encryption key, which wireless devices use to encrypt and decrypt wireless traffic. It's worth stressing that the PSK isn't the encryption key itself, but instead the PSK is used to create the key.



TIP This encryption key is used by either AES or TKIP for symmetric encryption depending on whether the system is using WPA2 or WPA. Chapter 14 explores symmetric encryption and AES in more depth.

Administrators enter the PSK into the WAP or wireless router, and users then enter the same PSK into their wireless devices. For example, Figure 3-16 shows one of the setup pages for a Linksys wireless router. The key is IW!!!Be\$\$\$PCertified. You can also see that WPA2-Personal is selected, and on this router, WPA2-Personal uses AES-based CCMP.

WPA-Enterprise and WPA2-Enterprise

An organization can increase security with wireless networks by adding an 802.1x authentication server and using Enterprise mode. A Remote Authentication Dial-In User Service (RADIUS) is commonly implemented as the 802.1x authentication server. Both WPA and WPA2 support Enterprise mode.

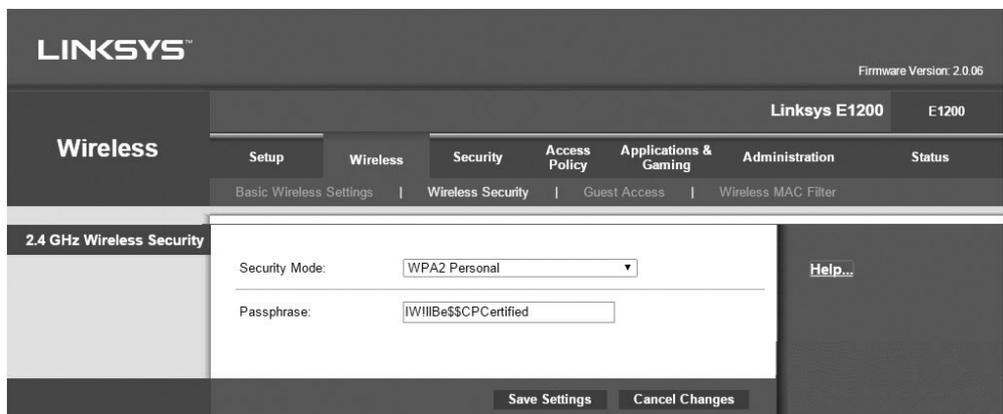


Figure 3-16 Configuring security for a wireless router



EXAM TIP Enterprise mode requires an 802.1x authentication server and increases security by requiring clients to authenticate before granting them access.

When an 802.1x authentication server is added, the clients are required to authenticate (such as with a username and password) prior to being granted access to the network. Additionally, clients don't need to enter a PSK. Instead, the system creates and assigns unique encryption keys after the clients log on. In contrast, Personal mode requires users to know and enter the PSK shared by all users.

An 802.1x authentication server will often be configured to check an authentication database on another server. For example, in a Microsoft domain, the 802.1x authentication server can check a user's credentials via a domain controller hosting Active Directory Domain Services. Users log on using their domain credentials. If they don't have an Active Directory account, they cannot log on and cannot access the wireless network.

Although the focus in this section is to use an 802.1x authentication server to authenticate wireless clients, an 802.1x authentication server can also be used to authenticate other clients. For example, it can be used to authenticate DHCP clients prior to providing TCP/IP configuration information.



EXAM TIP WEP should not be used because it has several security vulnerabilities. At the very least, WPA should be used. WPA2 is better, but if legacy hardware doesn't support WPA2, WPA might be the only alternative to WEP. WPA-Personal and WPA2-Personal use a PSK. WPA-Enterprise and WPA2-Enterprise use an 802.1x authentication server, providing a higher level of security.

Wireless Device Administrator Password

Wireless devices commonly have a web-based administration program that you can access with a basic web browser. You can access this program using the IP address assigned to the router (often either 192.168.0.1 or 192.168.1.1). You enter the IP address as the URL in your browser. When you connect, the wireless router prompts you to provide the name of the administrator account and the administrator password.

You can check the manual (or the Internet) for the default name and password if necessary. A common default combination is **admin** (for the administrator account) and **admin** (as the password), although there are many variations.

One of the important steps in securing a wireless device is to change this password from the default. Ideally, you'll change it to a strong password or passphrase. If you don't, an attacker might see your wireless network, log in using the defaults, and modify your network. An attacker can actually lock you completely out of your network until you reset the wireless device to factory defaults.



TIP Wardriving is the practice of driving around with a wireless receiver (such as a laptop with a wireless NIC) and looking for wireless networks.

Wireless Service Set Identifier

The service set identifier (SSID) is the name of the wireless network. Many wireless devices have default SSIDs, but newer ones often require the user to enter the name of the SSID as part of the setup process. For example, some common SSIDs for older wireless routers were **Linksys** (for Linksys systems) and **Netgear** (for Netgear devices).

An SSID can be almost any name you want, although it can't be longer than 32 characters. It's a user-friendly name, but you don't necessarily want to make it user friendly to attackers. For example, if your wireless network includes secret financial data, you probably don't want to call it **SecretFinancialData**.

Figure 3-17 shows the configuration page for a Cisco wireless router with the SSID entered as **SSCPWireless**. You can also see that the Broadcast SSID setting is checked, enabling SSID broadcast. When the wireless device is set to enable wireless SSID broadcasting, it regularly broadcasts messages with the SSID name. This enables other wireless devices in range of the broadcast to easily identify the wired network, and if they have security configured to match the wireless router, the SSID broadcast allows these other devices to connect easily.



TIP Figure 3-17 shows other settings. It's possible to convert this to a wired router only by setting the Wireless Interface to Disable. The Network Mode is set to N Only using 2.4 GHz, so it only supports 802.11n devices on 2.4 GHz. It can also be configured to support 802.11b, 802.11g, and 802.11n devices in mixed mode. This page also allows you to select the wireless channel, which is useful to avoid interference on another channel.

The screenshot displays the configuration interface for a Cisco DPC3825 DOCSIS 3.0 Gateway. The main navigation bar includes Setup, Wireless, Security, Access Restrictions, Applications & Gaming, Administration, Status, and Log OFF. The sub-navigation bar shows WPS, Radio Settings, Wireless Security, MAC Filter, Advanced Settings, WDS Settings, and QoS. The current page is titled "Wireless Network:" and contains the following settings:

- Wireless Interface: Enable Disable
- Network Mode: N Only (with a Scan 2.4GHz APs button)
- Standard Channel: 11 (Current: 11, Interference Level: Acceptable)
- Wireless Network Name (SSID): SSCPWireless
- BSSID: 7C:05:07:CB:C9:4B
- Broadcast SSID:

Buttons for "Save Settings" and "Cancel Changes" are located at the bottom of the configuration area.

Figure 3-17 Configuring the SSID

An ongoing discussion related to the SSID is whether you should enable or disable SSID broadcasting. Many people originally recommended disabling SSID broadcasting to hide the network. However, this really doesn't provide security. Even though this stops the normal SSID broadcasts used by other wireless devices, it doesn't prevent the device from actually broadcasting its SSID. An attacker with a wireless sniffer can easily capture key packets and discover the SSID, even when SSID broadcasting is disabled.

MAC Filtering

Most wireless devices also provide the ability to filter connections based on MAC addresses. For example, Figure 3-18 shows a configuration page for a wireless router with a Wireless MAC Filter set to Enable. As configured, it will permit one computer with the MAC address of 1A:2B:3C:4D:5E:6F access to the network. It will block access to all other computers.

However, it's quite simple to spoof the MAC address of a system. An attacker with a wireless sniffer can identify the MAC addresses that are connecting to a wireless network. Armed with this knowledge, the attacker can then use software to change the wireless MAC address of the attacker's system and then connect.

In theory, MAC filtering as a security precaution sounds good. In practice, it's easily defeated. The best protection you can use is WPA2. If you have a choice and can configure an 802.1x authentication server, use WPA2-Enterprise instead of WPA2-Personal (WPA2-PSK).



Figure 3-18 Enabling a MAC filter

Bluetooth

Bluetooth is a wireless technology used for small wireless networks commonly called personal area networks (PANs). As an example, smaller wireless devices such as smartphones along with an earpiece and microphone use Bluetooth. A user can wear a Bluetooth-enabled earpiece while keeping a Bluetooth phone in their pocket.

Devices have to be paired with each other before they will communicate. Pairing is done by changing the devices to discovery mode so that they can easily locate each other and connect. Once they are paired, discovery mode should be disabled. If discovery mode is left enabled, an attacker can pair with the devices and take them over. It's relatively easy for a user running Linux to exploit a Bluetooth device left in discovery mode.

The following are three well-known types of Bluetooth attacks:

- **Bluesnarfing** A hacker that is able to pair with a Bluetooth-enabled smartphone can access data on the device. The attacker can access information on the user's phone such as e-mail, contact lists, calendars, and text messages. The attacker must be close to the phone (usually within 30 feet).
- **Bluebugging** In some instances, an attacker is able to access a smartphone and issue specific commands to it. This enables the attacker to make phone calls, enable call forwarding, send messages, read and write data, connect to the Internet, and eavesdrop on conversations.
- **Bluejacking** Compared to Bluesnarfing and Bluebugging, Bluejacking is relatively harmless. It allows users to send unsolicited messages to nearby Bluetooth devices, but it does not involve modifying a phone's data.

GSM

Global System for Mobile Communications (GSM) began as a European standard for 2G, or digital cellular, networks in 1982. As a comparison, 1G refers to the analog communications that came out in the 1980s, and 2G refers to the digital technologies that came out in the 1990s. Industries embraced the GSM standard and it has since grown into a large family of technologies used in mobile communications.

GSM is a standard used in 219 countries (including the United States). The GSM Association (GSMA) includes almost 800 mobile operators and 200 other companies that help standardize, deploy, and promote the GSM standard.

3G, LTE, and 4G

3G refers to the third generation of standards for wireless access used by mobile phones and mobile telecommunications. 3G provides higher transmission speeds for mobile communications, comparable to broadband speeds that wired users have. Speeds up to 3 Mbps are common. Users can surf the Web with a mobile phone and access other web-based services such as e-mail.

Improvements of 3G are sometimes dubbed 3.5G and 3.75G. They use High Speed Packet Access (HSPA), and in many cases, users can get 20 Mbps download times or better.

However, not everyone can achieve these high rates of downloads. Worse, when large numbers of cellular users are operating in the same area, all users have reduced speeds because they're competing for the same bandwidth.

Telecommunication companies often market Long-Term Evolution (LTE) as 4G, but it doesn't always match the speed of 4G or meet the technical specifications. 4G (the fourth generation of mobile telecommunications) includes speeds as high as 100 Mbps for highly mobile users in trains and cars, and 1 Gbps for individuals walking or staying still.

New generations of telecommunications appear about every ten years. Fully compliant 4G systems were standardized in 2012. With that in mind, 5G systems are likely to appear sometime after 2020.

WiMAX

Worldwide Interoperability for Microwave Access (WiMAX) is a standard based on IEEE 802.16. IEEE 802.16 is officially known as Wireless MAN, but is commercially known as WiMAX. WiMAX provides broadband wireless coverage to large areas or metropolitan area networks (MANs).

For example, a WiMAX fixed station can provide wireless access for up to 30 miles. In comparison, wireless local networks using 802.11 are limited to about 300 feet. Many providers have created WiMAX networks around the world, and users can rent access to the networks. WiMAX provides speeds up to 40 Mbps for wireless connections. Future versions are expected to offer speeds up to 1 GBps.

Radio Frequency Identification

Radio frequency identification (RFID) is used for identification, tracking, asset management, and inventory control. Products are “tagged” with an RFID tag, which is a small electronic device that marks the product. RFID readers can then read the tag to get information about the product.

RFID tags can be either active (with a battery) or passive (without a battery). Active RFID tags are constantly transmitting. A passive RFID tag works similarly to how proximity cards work for authentication. An RFID reader can excite the electronics on the RFID tag and cause the tag to transmit its information.



TIP Chapter 2 describes proximity cards in more depth within the “Something You Have” section.

NFC

Near field communication (NFC) is a technology available on many smartphones that allows users to make purchases with those smartphones. For example, if you have an NFC-enabled smartphone, you can make a purchase at a store by simply tapping your smartphone on the NFC reader or waving it close to the reader. Similarly, users can share

information with each other by identifying the information to share and tapping the two devices together.

NFC uses similar technology as a proximity card, allowing devices to communicate with each other when they are in close proximity. The devices transmit and receive data only between short distances (identified as a *near field*). Additionally, NFC encrypts this data, providing some protection.

Of course, if you add your credit information to your NFC-enabled smartphone, it does make the risk of losing your phone that much greater. If someone steals your phone and can access it, they can wave and tap it during a massive shopping spree that will be billed on your credit card.

Protecting Mobile Devices

The primary vulnerabilities with mobile devices occur if they are lost. If someone finds or steals a mobile device, that person may be able get some valuable information. With this in mind, users can do several things to protect their devices. These include the following:

- **Password protect devices** Most devices support password protection. Users must enter a password or personal identification number (PIN) to unlock the device to use it. If a thief gets the device, this password makes it more difficult for the thief to access the data.
- **Encrypt data** Many devices support data encryption, and some even support full device encryption. If a thief does get an encrypted smartphone, the encryption provides confidentiality by restricting the thief's access to the data.
- **Remote wipe** Some devices support remote wipe. This allows the owner to send a command to a lost device to delete all the data on the device. The command will delete cached data, such as cached passwords, along with other data such as information in the user's contact list. For mobile devices that do not have a built-in remote wipe capability, users can install apps to perform the same function.
- **Enable GPS** Many devices support global positioning system (GPS) location services. If the device is lost, the user can use the GPS service to determine its location.



EXAM TIP Mobile devices are easily lost. Because they can often contain valuable data, it's important to know the different methods available to protect data. Password protection and data encryption are valuable if done before the device is lost. Remote wipe and GPS are valuable after the device is lost.

Chapter Review

The OSI Model is a framework for connecting multiple computers to each other. The ISO developed it and it has seven layers. The layers, listed from layer 1 to layer 7, are Physical, Data Link, Network, Transport, Session, Presentation, and Application. The TCP/IP Model (also called the TCP Model and the DoD Model) is similar to the OSI Model, but it has fewer layers. RFCs 1122 and 1123 identify four layers of the OSI Model, but some references in the SSCP CIB identify five layers.

Two important protocols operate on the Transport layer (layer 4). They are TCP and UDP. TCP is connection oriented and establishes sessions with a three-way handshake. UDP is connectionless and doesn't use a three-way handshake to establish a session. Instead, UDP uses its best effort to send the data. Other protocols can work with either TCP or UDP, depending on the protocol's needs. For example, FTP uses TCP for guaranteed delivery, while TFTP uses UDP.

Network configurations include bus, star, tree, token ring, and mesh. In a bus configuration, all devices share a common bus cable connection that has terminators at both ends. A bus is inexpensive to install but very difficult to troubleshoot, increasing the total cost of ownership. A star configuration is much more common and uses a central device such as a switch. All devices connect to each other via the central device in a star. A tree is a combination of bus and star topologies. A token ring configuration uses a logical token, and devices in the ring can transmit only when they have the token. Traditional token ring configurations are slower and don't scale well. FDDI uses a token ring concept but adds a second ring for redundancy and has significantly higher speeds because it uses fiber-optic cables. Mesh networks have multiple redundant connections, connecting each device with all other devices, and they provide the highest availability.

TCP/IP includes a full suite of protocols used by computers on the Internet and on many internal networks. The SSCP exam doesn't expect you to be an expert on all the protocols within TCP/IP, but you should be familiar with many of them.

DHCP provides IP addresses and other TCP/IP configuration information to hosts. ARP resolves IPv4 addresses to MAC addresses and must be running to transfer data on a network. RARP is similar to DHCP in that it provides IP addresses to clients with MAC addresses (also known as physical or hardware addresses). BootP provides diskless clients with an operating system. NDP is similar to ARP but resolves IPv6 addresses to EUI-64 addresses. DNS uses a hierarchical naming system with a distributed database to resolve host names to IP addresses.

ICMP is used for diagnostics. While it is useful when it is running, it is often disabled or blocked at firewalls to thwart attacks that use ICMP. IPv4 uses IGMP for multicasting, while IPv6 uses ICMPv6 for multicasting. SNMP is used to manage network devices. Agents receive data on UDP port 161 and send traps and other data over UDP port 162.

FTP is efficient for sending large files, while TFTP is good for sending small files such as configuration files for network devices. FTP uses TCP ports 20 and 21, and supports authentication. TFTP uses UDP port 69 and does not include authentication. Both FTP and TFTP send data across a network in cleartext.

Sending data over a network in cleartext is a security risk because an attacker with a sniffer can capture these transmissions and read the data. Several protocols, such as Telnet

and rlogin, send credentials over a network in cleartext. SSH is a more secure alternative as it encrypts the transmissions using a combination of symmetric and asymmetric encryption for confidentiality. SSH uses hashing for integrity, and can use an authenticated tunnel for mutual authentication. SSH can also encrypt other protocols such as FTP (as SFTP) and Secure Copy. SSH uses TCP port 22.

HTTP is the common protocol used to transfer web pages over the Internet. HTTPS encrypts e-commerce transactions with SSL or TLS. TLS is the designated replacement for SSL. HTTP uses TCP port 80 and HTTPS uses TCP port 443.

Routers use routing protocols such as RIPv2 and OSPF to share information with each other. RIPv2 is good for smaller networks, while OSPF works better in larger networks. Both protocols operate on the Network layer.

E-mail protocols include SMTP, POP3, and IMAP4. Clients send e-mail with SMTP and receive e-mail with POP3. IMAP4 allows clients to access and manipulate e-mail on a server, including organizing e-mail within folders. SMTP uses TCP port 25, POP3 uses TCP port 110, and IMAP4 uses TCP port 143.

VPNs use tunneling protocols to protect data. PPTP and L2TP are two tunneling protocols, and both operate on the Data Link layer. PPTP provides encryption, while L2TP depends on IPsec for encryption. IPsec uses AH for authentication and integrity. IPsec ESP includes the functionality of AH and provides confidentiality by encrypting the data. AH uses protocol number 51, and ESP uses protocol number 50.

Both TCP and UDP can use any of 65,536 ports (numbered from 0 to 65,535). Ports 0 to 1023 are well-known ports, and IANA maps them to specific protocols. For example, SSH uses port 22, HTTP uses port 80, and SMTP uses port 25. You should be able to identify the port used by many common protocols, in addition to the protocol numbers used by other protocols.

Different network architectures provide different levels of trust. For example, an intranet is internal within an organization and provides a high level of trust. The Internet is external and has the least level of trust. A DMZ hosts Internet-facing servers that are accessible by anyone on the Internet, and an extranet hosts Internet-facing servers that are accessible only by trusted entities such as business partners. Internal clients use private IP addresses, and Internet hosts have public IP addresses. NAT translates IP addresses from public to private and from private back to public.

Wireless networks allow clients to connect using radio-wave transmissions. These radio-wave transmissions are the easiest to capture and intercept (compared to transmissions on wired networks), but some basic security steps can increase the security of a wireless network. WEP is older and should not be used. WPA was an interim replacement for WEP and was compatible with existing hardware. WPA2 is the permanent replacement and is formally defined by 802.11i. WPA originally used TKIP, but many WPA systems also support AES. WPA2 uses AES-based CCMP. Both WPA and WPA2 can use Personal mode (with a preshared key) or Enterprise mode. Enterprise mode requires an 802.1x authentication server.

Smaller wireless devices use technologies such as GSM, 3G, LTE, 4G, WiMAX, and NFC. Several methods protect these devices if they are lost. Methods include password-protecting the devices, encrypting data on them, using remote wipe to delete data on lost devices, and enabling GPS to locate lost devices.

Questions

1. Which layer of the OSI Model defines cable standards?
 - A. Physical layer
 - B. Data Link layer
 - C. Network layer
 - D. Transport layer
2. Which layer of the OSI Model packages data as a frame?
 - A. Physical layer
 - B. Data Link layer
 - C. Network layer
 - D. Transport layer
3. Which layer of the OSI Model handles physical addressing?
 - A. Physical layer
 - B. Network layer
 - C. Data Link layer
 - D. Transport layer
4. Which layer of the OSI Model packages data as a packet?
 - A. Physical layer
 - B. Data Link layer
 - C. Network layer
 - D. Transport layer
5. Which layer of the OSI Model provides reliable end-to-end communication services?
 - A. Physical layer
 - B. Transport layer
 - C. Data Link layer
 - D. Host layer
6. Which layer of the OSI Model includes TCP and UDP?
 - A. Transport layer
 - B. Network layer
 - C. Data Link layer
 - D. Application layer

7. Which of the following protocols is connection oriented?
 - A. IP
 - B. RIP
 - C. TCP
 - D. UDP
8. Which layer of the TCP/IP Model corresponds to the OSI Network layer?
 - A. Host layer
 - B. Application layer
 - C. Internet layer
 - D. Link layer
9. Which of the following topologies avoids collisions using a token?
 - A. IEEE 802.3
 - B. IEEE 802.5
 - C. CSMA/CD
 - D. CSMA/CA
10. What protocol would a system use to determine a system's physical address?
 - A. ARP
 - B. RARP
 - C. BootP
 - D. DNS
11. Which of these ports does DNS use?
 - A. TCP 23
 - B. TCP 25
 - C. UDP 53
 - D. UDP 69
12. Which of the following protocols is commonly used with diagnostic utilities?
 - A. TFTP
 - B. RARP
 - C. IGMP
 - D. ICMP

13. Which of the following accurately identifies a difference between FTP and TFTP?
 - A. FTP uses UDP, and TFTP uses TCP.
 - B. FTP supports authentication, but TFTP does not support authentication.
 - C. TFTP sends data across a network in cleartext, but FTP encrypts data.
 - D. TFTP is primarily used to transfer large files, and FTP is used to transfer configuration information to and from network devices.
14. Which of the following protocols is a more secure alternative for remote login?
 - A. Telnet
 - B. rlogin
 - C. rexec
 - D. SSH
15. What port does POP3 use?
 - A. 25
 - B. 110
 - C. 143
 - D. 443
16. You are purchasing a product from a website. Which of the following protocols will your system most likely use to provide confidentiality for this transaction?
 - A. SSL
 - B. SSH
 - C. IPsec
 - D. HTTP
17. Which of the following statements is correct related to IPsec?
 - A. IPsec provides confidentiality by encrypting data with AH.
 - B. IPsec provides confidentiality by encrypting data on the Network layer.
 - C. IPsec AH uses protocol number 50.
 - D. IPsec ESP uses protocol number 51.
18. What is the protocol number for IPsec AH?
 - A. 1
 - B. 6
 - C. 50
 - D. 51

19. Where is a DMZ located?
- A. Behind the intranet firewall
 - B. In front of the first intranet-facing firewall
 - C. In front of the first Internet-facing firewall
 - D. Behind the first Internet-facing firewall
20. Which of the following is the recommended security mechanism to use with wireless networks?
- A. 802.11a
 - B. 802.11g
 - C. 802.11i
 - D. 802.11n

Answers

1. **A.** Cable standards are defined at the Physical layer, layer 1. They are not defined at the Data Link layer (layer 2), the Network layer (layer 3), or the Transport layer (layer 4).
2. **B.** The Data Link layer packages data as a frame. The Physical layer packages data as bits. The Network layer packages data as a packet. The Transport layer packages data as a segment.
3. **C.** The Data Link layer uses physical addresses, also called hardware addresses and media access control (MAC) addresses. The Physical layer packages data as bits and doesn't use addresses. The Network layer uses IP addresses (also called logical addresses). The Transport layer doesn't use addresses but uses ports to identify traffic.
4. **C.** The Network layer packages data as a packet. The Physical layer packages data as bits. The Data Link layer packages data as a frame. The Transport layer packages data as a segment.
5. **B.** The Transport layer provides reliable end-to-end communication services. Neither the Physical layer nor the Data Link layer provides this service. The Host layer is on the TCP/IP Model, not the OSI Model.
6. **A.** The Transport layer includes the TCP and UDP protocols. These protocols are not implemented on the Network layer, the Data Link layer, or the Application layer.
7. **C.** TCP is connection oriented. IP uses TCP to provide a connection-oriented session but is not connection oriented itself. RIP is a routing protocol and is not connection oriented. UDP is connectionless. Instead of establishing a session, it makes a best effort to deliver data.

8. **C.** The TCP/IP Internet layer corresponds to the OSI Network layer. The TCP/IP Host (or Host-to-Host) layer corresponds to the OSI Transport layer. The TCP/IP Application layer corresponds to the Application, Presentation, and Session OSI layers. The TCP/IP Link layer (also called the Network Interface or Network Access layer) corresponds to the OSI Data Link and Physical layers.
9. **B.** IEEE 802.5 defines token ring networks, which avoid collisions using a token. Ethernet (IEEE 802.3) attempts to detect collisions using Carrier Sense Multiple Access with Collision Detection (CSMA/CD). Wireless networks (802.11) attempt to avoid collisions using Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA).
10. **A.** Systems use the Address Resolution Protocol (ARP) to identify the assigned physical (or MAC) address matching an assigned IP address. Reverse ARP (RARP) allows a system with a MAC address to get an IP address. The Bootstrap Protocol (BootP) allows a diskless system to get an IP address and then download the image of an operating system. Domain Name System (DNS) resolves host names to IP addresses but not physical addresses.
11. **C.** Domain Name System (DNS) uses UDP port 53 when clients query the DNS server and TCP port 53 when DNS servers transfer data between each other. Telnet uses TCP port 23. Simple Mail Transfer Protocol (SMTP) uses TCP port 25. Trivial FTP (TFTP) uses UDP port 69.
12. **D.** Many diagnostic utilities such as ping, pathping, and tracert use Internet Control Message Protocol (ICMP). Administrators commonly use Trivial FTP (TFTP) to transfer configuration files to and from network devices. Reverse ARP (RARP) allows a system with a MAC address to get an IP address. The Internet Group Message Protocol (IGMP) is used for IPv4 multicasting.
13. **B.** File Transfer Protocol (FTP) supports authentication, but Trivial FTP (TFTP) does not support authentication. FTP uses TCP ports 20 and 21, while TFTP uses UDP port 69. Both FTP and TFTP send data across a network in cleartext, but it is possible to encrypt FTP with Secure Shell (as SFTP). TFTP is commonly used to transfer configuration files to and from network devices, and FTP is primarily used to transfer large files.
14. **D.** Secure Shell (SSH) encrypts data sent over a network and is the most secure method for remotely accessing systems of the given choices. Telnet, rlogin (which is remote login), and rexec (remote execute) all send data across a network in cleartext.
15. **B.** Post Office Protocol version 3 (POP3) uses TCP port 110. Simple Mail Transfer Protocol (SMTP) uses TCP port 25. Internet Message Access Protocol version 4 (IMAP4) uses TCP port 143. HyperText Transfer Protocol Secure (HTTPS) uses TCP port 443.
16. **A.** E-commerce transactions use HyperText Transfer Protocol Secure (HTTPS) for confidentiality, and Secure Sockets Layer (SSL) is one of the protocols used to encrypt HTTPS. While not one of the choices, Transport Layer Security (TLS) is

also commonly used to encrypt HTTPS. While Secure Shell (SSH) and Internet Protocol security (IPsec) both provide confidentiality with encryption, HTTPS doesn't use SSH or IPsec. HTTP sends data in cleartext, so it doesn't provide confidentiality.

17. **B.** Internet Protocol security (IPsec) provides confidentiality by encrypting data on the Network layer. Encapsulating Security Payload (ESP) provides confidentiality by encrypting data, but Authentication Header (AH) only provides authentication and integrity. AH uses protocol number 51, and ESP uses protocol number 50.
18. **D.** The protocol number for Internet Protocol security (IPsec) Authentication Header is 51. The protocol number for IPsec Encapsulating Security Protocol (ESP) is 50. The protocol number for Internet Control Message protocol (ICMP) is 1, and the protocol number for Transmission Control Protocol (TCP) is 6.
19. **D.** A demilitarized zone (DMZ), or perimeter network, is located behind the first Internet-facing firewall. It is not on the private network (behind the intranet firewall) or directly on the Internet (in front of the intranet or Internet-facing firewall).
20. **C.** The 802.11i standard documents Wi-Fi Protected Access 2 (WPA2), the recommended security mechanism for wireless networks. It uses AES-based CCMP for very strong security. The other standards focus on the base frequency and speed of wireless networks, not security.

Advanced Networking and Communications

In this chapter, you will learn about

- The difference between switches and routers
- Telecommunications methods used to access the Internet
- Securing the Voice over Internet Protocol (VoIP) with the Secure Real-time Transport Protocol (SRTP)
- Packet-filtering, stateful inspection, and application firewalls
- Defense diversity with firewalls
- The differences between network-based and host-based firewalls
- Risks and vulnerabilities related to remote access solutions
- Different tunneling protocols used with remote access
- Authentication methods used with remote access
- Network access control
- The use of virtualization
- Cloud service models and operation models

Managing LAN-Based Security

A local area network (LAN) connects computing devices together, allowing users to share resources such as data and devices. While LANs have helped improve productivity for many people, they have also added many risks. Because of this, it's important for IT professionals to understand networks and some basic LAN-based security methods.

Comparing Switches and Routers

Switches connect devices together in a network and routers connect networks together. Figure 4-1 shows devices connected together via a switch in Network 1, devices connected together via a switch in Network 2, and both networks connected together via a router. Notice that Port 1 of the router is the gateway for computers in Network 1 to other networks. Similarly, Port 2 is the gateway for computers in Network 2 to other networks. Most networks have only one gateway labeled as the default gateway.



NOTE Networks connected together in the same area, such as within an office building or school, are local area networks (LANs). However, technicians often just refer to a LAN as a network.

Chapter 3 covers the Open Systems Interconnection (OSI) Model and mentions how switches operate on layer 2 (the Data Link layer) and routers operate on layer 3 (the Network layer). Layer 2 uses media access control (MAC) addresses, and switches control traffic within a network using the MAC addresses. Layer 3 uses Internet Protocol (IP) addresses, and routers control traffic between networks using IP addresses.

Switches typically build tables mapping physical ports to the MAC address of the device connected to the port. For example, when Computer 1 sends traffic through the switch, the switch captures its MAC address and records it in the table. While administrators can statically configure the MAC address table, most switches create the table dynamically. Table 4-1 shows a table for the switch in Network 1. Note that Computer 1 is connected to physical Port 2 and Computer 2 is connected to physical Port 4.

When Computer 1 sends unicast traffic to Computer 2, it includes the destination MAC address (34-64-a9-13-44-a8). The switch uses the MAC address table and sees that the destination MAC address is connected to Port 4, so it forwards the traffic to Port 4. This traffic doesn't reach any of the other ports.

In contrast, older networks used hubs in place of switches. A hub forwards traffic going into one port to all other ports. This represents a risk because an attacker could configure a protocol analyzer (or sniffer) to capture traffic going through the hub. Chapter 5 covers sniffing attacks in more depth.



NOTE Layer 3 switches are an advanced implementation of switches. They work similarly to how routers work on layer 3.

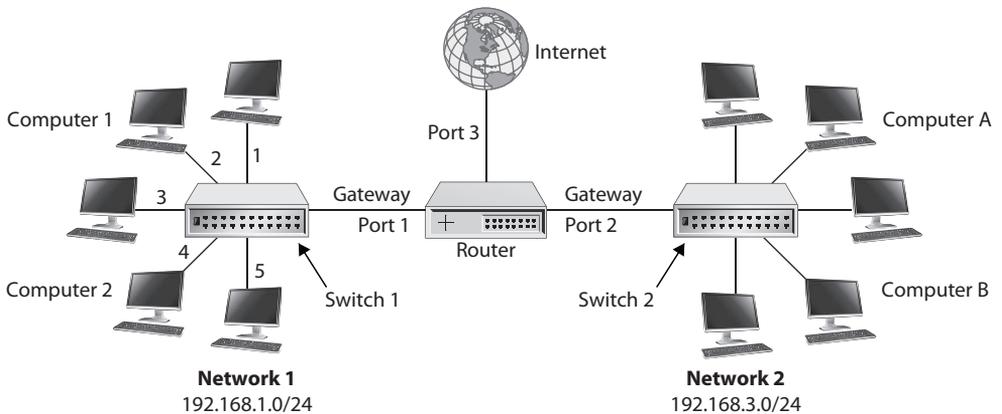


Figure 4-1 Comparing switches and routers

Table 4-1
MAC Address
Table in a Switch

Physical Port	MAC Address
1	76-29-af-e4-33-99
2 (Computer 1)	74-29-af-e4-33-9a
3	34-64-a9-13-44-08
4 (Computer 2)	34-64-a9-13-44-a8
5	00-14-22-a3-32-c8

Routers use tables too, but they use IP addresses and network addresses instead of MAC addresses. As an example, the routing table for the router in Figure 4-1 would include the following information:

- Physical Port 1 is connected to Network 1 (all IP addresses in the 192.168.1.0/24 network).
- Physical Port 2 is connected to Network 2 (all IP addresses in the 192.168.3.0/24 network).
- The path to all unknown networks (such as those on the Internet) is through Port 3.



NOTE The /24 in 192.168.1.0/24 indicates the subnet mask is 255.255.255.0. It is using Classless Inter-Domain Routing (CIDR) notation and indicates the first 24 bits in the subnet mask are 1's.

Administrators can create routing tables manually or configure routers with routing protocols. Routing protocols such as Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) allow routers to share information with other routers and learn the best path to a different network.

Segmentation

Network segmentation refers to breaking up a larger network into smaller interconnected networks. Figure 4-1 shows how a router segments traffic between two networks. The router divides a large network into two smaller networks. Another method of segmenting a network is with virtual local area networks (VLANs). A VLAN allows an organization to segment traffic with a lot more flexibility when compared to segmenting traffic with a router.

Consider the single-network configuration shown in Figure 4-2. Imagine that Switch 1 connects the computers on a building's first floor, and Switch 2 connects the computers on a building's second floor. Switch 3 is connecting the two switches to the router. Configured like this, it creates a single broadcast domain. In other words, if any computer transmits broadcast traffic, it will reach all of the computers on the first and second floors. Routers do not pass broadcast traffic, so the broadcast domain does not extend beyond the router.

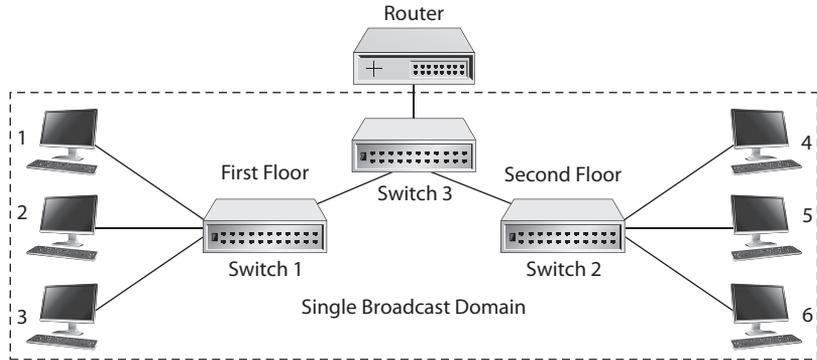


Figure 4-2 Building network on two floors

While this configuration works, it also presents a potential risk. Imagine that employees using computers 1 and 4 are working on a sensitive project. Data broadcasts between these two computers will go to all computers in the broadcast domain, meaning someone within the broadcast domain can capture this traffic with a sniffer.

One way to segment these two computers is to add another switch on the other side of the router. You can then connect the two computers to the switch as shown in Figure 4-3. However, this is highly inefficient. You have to add a switch and run additional cables from the switch to the computers on two separate floors.

VLANs provide a much more efficient solution, as shown in Figure 4-4. Computers 1 and 4 are within the VLAN named VLAN 10. This creates a logical broadcast domain for these two computers just as if they were configured with another switch as shown in Figure 4-3. Broadcast traffic in VLAN 10 will not reach any of the devices on the

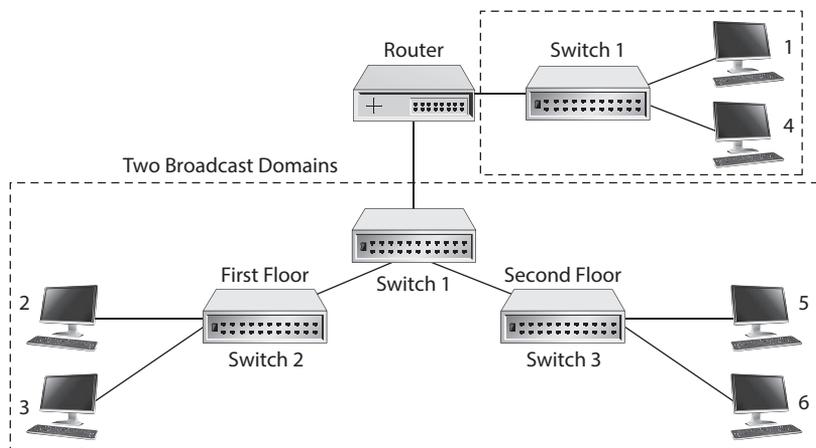


Figure 4-3 Inefficient network segmentation

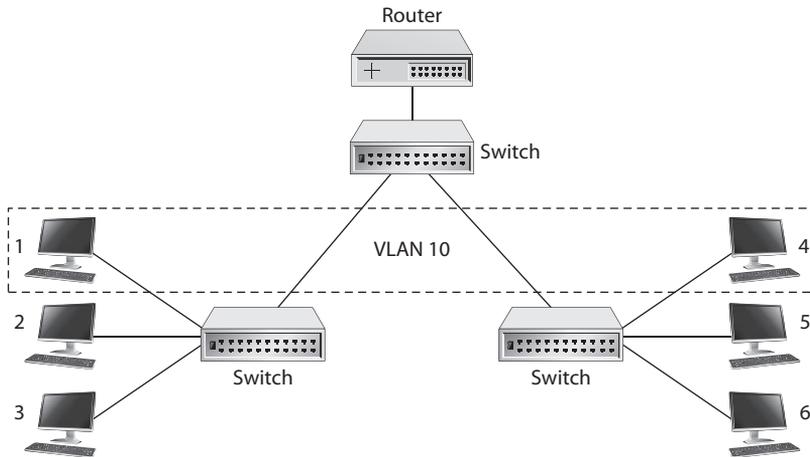


Figure 4-4 VLAN configuration

network. Similarly, broadcast traffic outside of VLAN 10 will not reach computers 1 and 4. Less broadcast traffic also improves the performance of the LAN.



NOTE When configuring VLANs, it's important to configure all of the switches within the network to recognize the VLAN. This ensures that broadcast traffic within a VLAN does not reach devices connected to other switches.

While Figure 4-4 shows a single VLAN, it's also possible to create multiple VLANs. For example, you could create a VLAN for computers 2 and 5, and another VLAN for computers 3 and 6. VLANs provide several benefits:

- **Security** VLANs separate traffic from one broadcast domain to another. This is especially useful if sensitive data is transmitted over a network.
- **Improved performance** VLANs reduce the size of broadcast domains and improve their performance because there is less traffic within each broadcast domain.
- **Reduced costs** A VLAN doesn't require any additional hardware as long as the switches support VLANs.



EXAM TIP VLANs increase security by segmenting traffic on a network. Administrators create VLANs with switches.

Secure Device Management

It's important to protect routers and switches with strong physical security. A data center or server room already has strong physical security, so it is a good location for routers and switches. Another option that many organizations use is to store them in wiring closets. However, it's important to ensure that the wiring closet is kept locked to prevent unauthorized access.

If attackers can access the devices, they can hook up a sniffer and capture traffic going through them. As an example, an attacker can connect a rogue access point to a switch. The access point will then transmit captured data wirelessly. The attacker can collect the data using a laptop while sitting in the parking lot outside the organization.

By default, a switch only forwards broadcast traffic to other ports, so the attacker cannot see all traffic going through the switch. However, most switches (and some routers) support port mirroring. Port mirroring sends a copy of all packets going through a switch to a separate port used for monitoring. Cisco switches support port mirroring with a feature called Switched Port Analyzer (SPAN), or the newer Remote SPAN (RSPAN) and Encapsulated RSPAN (ERSPAN). If the attacker can configure port mirroring, it's possible to capture all traffic going through the switch.



NOTE While Cisco routers don't support port mirroring with SPAN, RSPAN, or ERSPAN, there are other methods to achieve the same result. For example, many Cisco routers support the IP Traffic Export feature, which will send a copy of all traffic to a monitored port. Similarly, Juniper routers and switches include methods for port mirroring with their devices.

Understanding Telecommunications

Telecommunications is the transmission of signals for communications. It includes connections with phones, copper and fiber-optic cables, satellites, and more. Data transmitted over any of these connections faces certain risks, such as the risk of capture. For example, an attacker can use a sniffer to capture traffic going over the Internet. Once captured, the attacker can view and analyze the data.

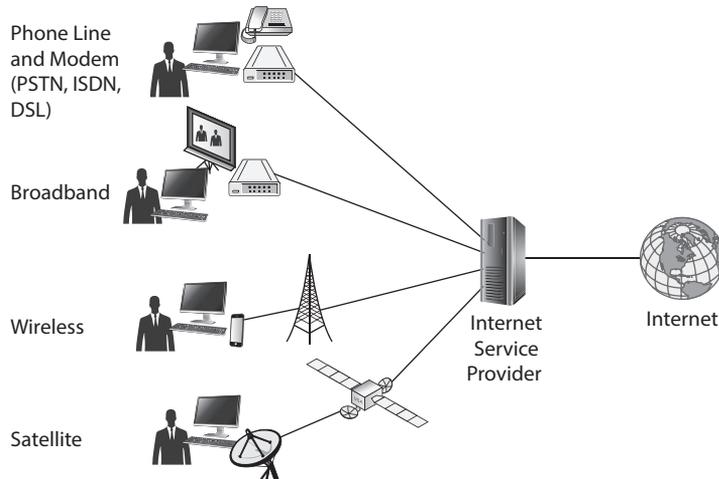
A significant amount of telecommunications traffic goes over the Internet. This includes typical Internet traffic such as web pages, e-mail, streaming audio and video, and some types of voice communications.

Internet Connections

People use multiple methods to connect to the Internet. Although there have been many changes over the years, many of the basic methods still exist. Figure 4-5 shows many of the methods that users employ to connect to the Internet via an Internet service provider (ISP), and the following list explains these methods:

- **PSTN** The public switched telephone network (PSTN) includes the wired and wireless infrastructure used to support the world's telephone system. Plain old telephone service (POTS) is the term for traditional phone lines that provide

Figure 4-5
Methods of
connecting to
the Internet



phone service to a significant portion of the world's population. Anyone with access to a telephone is able to dial in to the Internet through an ISP. The maximum speed of dial-up lines is typically only about 50 Kbps.

- **ISDN** An integrated services digital network (ISDN) line provides access to the Internet through a circuit-switched telephone network. ISDN lines are typically 128 Kbps.
- **DSL** A Digital Subscriber Line (DSL) is a group of different technologies used to transmit digital signals over the telephone network. It originally started as ISDN DSL (IDSL) and has been expanded into multiple versions. Technicians often refer to the different versions as xDSL and the most common version in use today is Asymmetric DSL (ADSL). DSL services can provide speeds starting at 256 Kbps with maximums as high as 40 Mbps.
- **Cable modems** Cable TV companies long ago realized that they could add Internet signals onto the existing cable and provide Internet access to cable TV customers. The cable already existed, so it was just a matter of adding some equipment to manage the Internet connections. Cable access is commonly called broadband Internet access, or just broadband, and often provides signal speeds as high as 4 Mbps.
- **Cellular** The coverage of cellular phone companies has been growing steadily over the years, and although they don't yet cover 100 percent of the United States, they do cover a substantial amount of the land. When the signal is data enabled, users can use a wireless cellular connection for Internet access. Many smartphones include a hotspot capability that allows other devices to connect to the Internet through the smartphone.
- **Satellite** Satellite Internet access is available in some rural areas. Users are able to use an ISP's equipment to upload and download Internet data via a satellite with a two-way satellite dish at their home.

Internet Connections in Rural Areas

Although the use of dial-up connections is almost zero in urban areas where broadband is widely available, dial-up is still used in rural areas where broadband cables are not available. According to a 2013 Pew survey, approximately 3 percent of Americans (about 9 million people) still use dial-up. A 2015 financial report from AOL indicates that nearly 2.1 million people are still signed up to their ISP service.

Satellite connections have been available for quite a while now, but customers often complain about inconsistent service from the satellite companies. A growing technology in urban areas within the United States is the use of wireless transmissions similar to what smartphone users can access for surfing the Internet. Although speeds vary depending on the strength of the connection, they are usually much better than that of a dial-up connection.

VoIP

Voice over Internet Protocol (VoIP) includes the different technologies used to transmit multimedia and voice communications over IP networks, including the Internet. It is possible to transmit VoIP over the Internet in a plain, unencrypted format, but such transmissions are easily intercepted. Instead, many VoIP implementations provide some type of protection for the transmissions.

RFC 3711 defines Secure Real-time Transport Protocol (SRTP), which provides confidentiality, message authentication, and replay protection for audio and video traffic, including VoIP. Chapter 14 covers cryptography topics in greater depth, but it's worth mentioning that SRTP supports strong cryptographic methods such as Advanced Encryption Standard (AES) for encryption, and salting to protect against offline attacks. Salting adds bits to passwords or encryption keys to thwart brute-force attacks.



EXAM TIP SRTP can protect VoIP transmissions and provides confidentiality, authentication, and replay protection.

Securing Phones

Many organizations use a private branch exchange (PBX) or a business telephone system to control the services provided to users in the organization. These systems include a wide range of functions, such as automated answering machine functions, call accounting (to track events such as long distance calling), call blocking, conference calling, caller identification, speed dialing, and more. Most systems also support remote administration of the system by telephone system administrators.

Some of the common methods used to secure phone systems include the following:

- **Protect the system and wiring with physical security** Store systems within a secure area such as in a server room and restrict access to only authorized individuals.

- **Protect and regularly change the administrator password** PBX systems have administrator accounts to access advanced functions. These passwords should be strong, kept private, and changed regularly.
- **Restrict numbers that can be used for call forwarding** Most PBX systems allow you to specify what in-house numbers will support call forwarding. Additionally, you can restrict the forwarding numbers. This prevents attackers from modifying call-forwarding features so that calls are forwarded to phones controlled by the attackers.
- **Restrict long distance calling** You can enable toll-free calling and block other long distance calling. Most systems include the ability to allow long distance calling in certain situations, such as when a caller enters an authorization number.



EXAM TIP Phone systems or PBXs should be protected with security. This includes ensuring the physical security of the system and controlling what users are able to do. Call forwarding should be blocked or at least highly restricted.

Converged Communications

Most of today's networks support standard protocols within the Transmission Control Protocol/Internet Protocol (TCP/IP) suite. However, these protocols don't always meet the needs of organizations. Instead, organizations sometimes integrate specialty or proprietary protocols into their networks. *Converged communications* refers to the merging of standard protocols with specialty or proprietary protocols.

Fibre Channel over Ethernet (FCoE) is an example of a converged protocol. It supports the use of Fibre Channel (FC) commands over a standard TCP/IP network. The original FC networks require specialized hardware and cabling, which is significantly more expensive when compared to using an existing TCP/IP network.

There are two security challenges with converged communications. First, the technologies aren't widely used, so technicians and administrators may not understand them fully. Because of this, they may not know the proper steps to configure them properly. Second, discovering vulnerabilities may take longer simply because not as many organizations are using the technologies.

Using Proxy Servers

A proxy server is used as an intermediary, or proxy, for clients. Instead of clients' retrieving web pages directly from the Internet, they send the request to the proxy server. The proxy server retrieves the web page on the client's behalf and returns it to the client.

Chapter 3 briefly introduced proxy servers in the context of Network Address Translation (NAT). Remember that NAT translates private IP addresses to public addresses and public IP addresses back to private. Proxy servers commonly include NAT as an additional service. Additionally, proxy servers provide two more primary services: caching and filtering. Figure 4-6 shows how caching and filtering work.

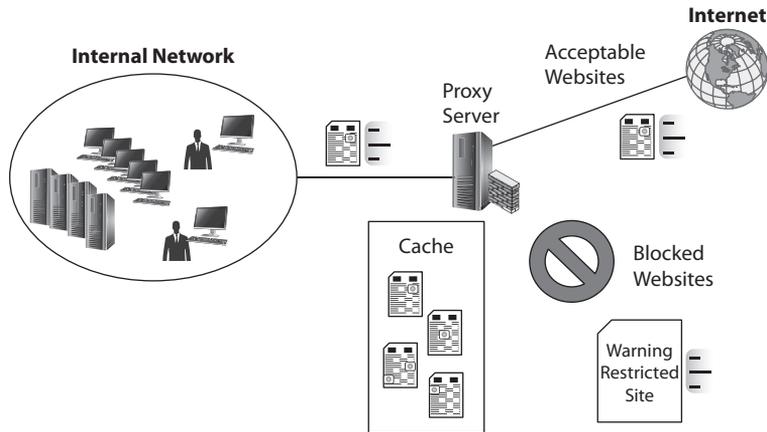


Figure 4-6 Proxy used for caching and filtering

When a user first retrieves a web page, the proxy server retrieves it from the Internet, stores a copy in cache, and sends a copy to the user. If another user requests the same page, the proxy server retrieves it from cache and sends the page to the other user instead of retrieving it from the Internet again. This improves performance for the clients and also reduces Internet bandwidth usage.



NOTE Many web pages include metadata that can indicate when a web page is no longer valid. If the web page times out based on this metadata, the proxy server flushes it from the cache. If a user requests the page again, the proxy server retrieves another copy from the Internet and places the new copy in cache.

Most proxy servers can filter traffic based on the web address. For example, if an organization wants to block employees from accessing websites with certain types of content, such as gambling sites, they can use a filter to block these websites. Many organizations purchase lists of websites in certain categories from third-party companies. For example, if an organization wants to block users from accessing web-based e-mail, it can purchase a list of websites in the web-based e-mail category.



TIP Lists of websites are available in multiple categories, including gambling, shopping, pornography, and many more. Organizations subscribe to the lists that they want to block, and the vendor updates these lists regularly as new sites appear.

Administrators install the list of websites on the proxy server, and when users attempt to access a site on the list, the proxy server blocks the request. Instead, the proxy server displays a page indicating that the proxy server blocked their query. Many organizations personalize this web page to remind users about the organization's security policy.

Some organizations have switched to whitelisting websites. For example, instead of blocking specific websites and allowing all others, an organization can list only the sites that are allowed and block all others. Although this does limit a user's ability to surf the Internet, that is sometimes the goal.



TIP A whitelist identifies sites that are authorized and a proxy server blocks all sites not on the whitelist. In contrast, a blacklist identifies unauthorized sites and a proxy server allows access to any other sites.

Understanding Firewalls

Chapter 3 introduced firewalls when describing differences between the Internet, an intranet, an extranet, and a demilitarized zone (DMZ). Network-based firewalls filter the traffic going into or out of a network. Host-based firewalls are software applications that run on individual systems such as a server or desktop computer. Firewalls have gone through several generations, as described in the following sections.

Packet-Filtering Firewall

The first generation of firewalls is packet-filtering firewalls that filter traffic based on the contents of a packet. More specifically, a packet-filtering firewall can filter traffic based on IP addresses, subnet addresses, ports, some protocols, or any combination of these:



EXAM TIP A packet-filtering firewall uses rules within access control lists (ACLs) to filter the traffic. ACLs filter traffic based on source or destination IP addresses, subnet addresses, entire domains, ports, and/or protocols.

- **IP addresses** Each packet contains the source and destination IP addresses. A packet filter can examine this data and either block or allow traffic to or from specific IP addresses. For example, if an attacker is launching an attack from a specific IP address, it's possible to create a rule to block traffic from that IP address.
- **Subnet addresses** In addition to filtering specific IP addresses, you can filter entire subnets. For example, the network ID of 192.168.1.0/24 could be used to block or allow IP addresses from 192.168.1.1 through 192.168.1.254. Similarly, rules can block entire blocks of IP addresses, such as those from another country.
- **Ports** Well-known ports (ports from 0 to 1023) identify many protocols. For example, the well-known port for the Simple Mail Transport Protocol (SMTP) is port 25. By allowing traffic through port 25 (or opening port 25), you are effectively allowing SMTP traffic.
- **Some protocols** Many protocols are identified by a protocol number (or protocol ID) embedded in the IP header. For example, the protocol number for Internet Control Message Protocol (ICMP) is 1. A rule that blocks traffic with a protocol number of 1 blocks all ICMP traffic.



TIP Chapter 3 covers ports in more depth and includes a table listing many commonly used ports. Additionally, Chapter 3 describes protocol numbers and includes a table listing many commonly used protocol numbers.

- **A combination of IP addresses, ports, or protocols** Any combination of IP addresses, ports, or protocols can be used. For example, consider Figure 4-7, which shows an e-mail server (named mail1) within a DMZ and another e-mail server (named mail2) in the intranet. Administrators would configure Firewall 1 to allow SMTP traffic (using port 25) to the mail1 server (using IP address 192.168.1.25). Similarly, administrators would configure Firewall 2 to allow e-mail traffic between mail1 and mail2. However, unless administrators add more rules to these firewalls, they will block SMTP traffic coming from the Internet and going directly to mail2.

A packet-filtering firewall uses an implicit deny philosophy. All traffic is blocked (implicitly denied) unless there is a rule in the ACL that explicitly allows the traffic. Some firewalls automatically apply this rule, while other firewalls require administrators to add it. The implicit deny rule is placed at the end of the ACL and might look like one of the following statements:

- deny any
- deny any any

The actual syntax is dependent on the router brand and model, but the most important point is that it is the last rule in the ACL. If a previous rule doesn't explicitly allow the traffic, this rule blocks the traffic. If you place this rule first in the ACL, the ACL will ignore any other rules in the ACL and deny all traffic.



TIP Most routers and packet-filtering firewalls automatically use a **deny any** or **deny any any** rule even if the rule hasn't been typed into the ACL. In other words, you may not see the rule in the ACL, but the firewall is enforcing the rule.

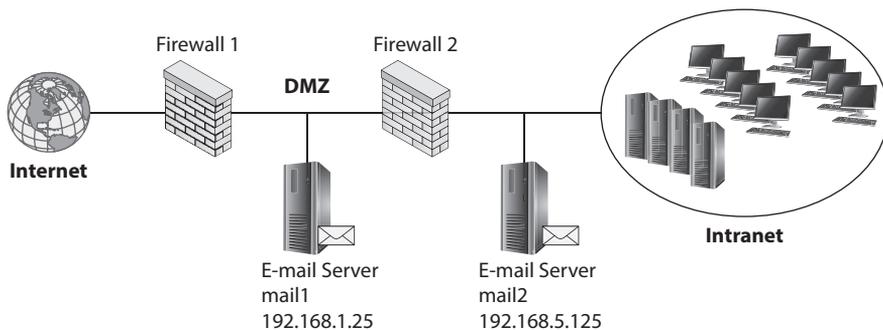


Figure 4-7 E-mail server within a DMZ

Although a packet-filtering firewall is effective at blocking individual packets, it can be difficult to maintain when the ACL includes a large number of rules. Additionally, a large number of rules can negatively affect performance because the firewall must compare each packet against each rule.

Many organizations use a simple packet-filtering firewall at their boundary to the Internet. They do so by adding rules to the ACL of an external screening router. Although such a firewall doesn't provide all the capabilities of a more advanced firewall appliance, it can be effective for smaller organizations.

Stateful Inspection Firewall

The second generation of firewalls is stateful inspection firewalls. They include the same capabilities of a packet-filtering firewall but can also filter traffic by examining much more than individual packets. Instead, they monitor the complete session. A stateful inspection firewall identifies active connections and monitors the status of these connections in a state table within the firewall.

As an example, a TCP session begins with the three-way handshake. The stateful inspection firewall uses this three-way handshake to identify that a session has started between two systems. The firewall can then monitor the traffic between the systems. Even though UDP connections don't start with a three-way handshake, a stateful inspection firewall can still identify active UDP connections and track the state of the UDP connection.



TIP A stateful inspection firewall allows traffic that is part of an active connection or that is initiating a new connection. It rejects traffic that is not part of an active connection or that is not initiating a new connection. It also includes basic packet-filtering capabilities of a packet-filtering firewall.

A stateful inspection firewall is aware of packets exchanged in many communication sessions. When a session starts, the firewall dynamically opens ports based on the needs of the connections, and when the session completes, it dynamically closes ports.

In comparison, the packet-filtering firewall is a stateless firewall. It cannot determine whether a packet is part of an ongoing session between two systems or is from an attacker trying to probe or attack a system.

Application Firewall

An application firewall (also called a third-generation firewall) includes different elements for different applications. For example, it would have a HyperText Transfer Protocol (HTTP) proxy to examine web traffic, an SMTP proxy to examine e-mail traffic, a File Transfer Protocol (FTP) proxy to examine FTP traffic, and so on. Each proxy has the ability to analyze specific traffic for the protocol that it is monitoring and understands the individual protocol commands. For example, an HTTP proxy can interpret Get and Put commands used to retrieve and write web page elements, and administrators can configure the proxy to allow or block specific commands.



TIP An application firewall is also known as an application proxy firewall and an application gateway firewall.

Although the application firewall does provide effective filtering for individual applications or protocols (such as HTTP, SMTP, and FTP), it requires a separate component for each individual protocol. If the firewall must filter many different protocols, this process can be extensive and negatively affect performance. However, if the firewall is being used only to filter a few protocols, it is manageable and can be highly efficient.

Application firewalls are often used with packet filtering. For example, the packet-filtering firewall can drop packets for protocols that aren't desired based on the port, while the application proxy firewall can perform deep inspection of traffic for protocols that are allowed.

Next-Generation Firewall

Next-generation firewalls integrate multiple capabilities into a single security appliance, and these are often referred to as a unified threat management (UTM) device. Additionally, vendors update the UTM devices to adapt to new threats as they emerge.

Some of the additional capabilities that are included in UTM devices include

- **Malware blocking** UTM appliances typically include antivirus capabilities to detect and block malware before it enters the network. This includes malware included as attachments or downloaded from websites. Many UTM appliances can often detect and block malicious scripts embedded in web pages.
- **Spam filtering** Some UTMs include spam filters to block unwanted e-mail. This can be especially useful because attackers commonly use e-mail to launch various attacks. Chapter 5 discusses many common e-mail attacks such as phishing, spear phishing, and whaling.
- **URL filtering** UTMs commonly include URL-filtering capabilities, similar to a proxy server. URL filters can block users from accessing external websites.

Defense Diversity

Chapter 1 introduced the concept of defense in depth. A solid security program includes multiple layers of security so that even if one layer breaks down, other layers still provide protection. Similarly, defense diversity can be used when implementing a DMZ by using firewalls from two separate vendors.

For example, consider Figure 4-8. The firewall separating the Internet from the DMZ is from Cisco, and the firewall separating the DMZ from the intranet is from Microsoft. Because the DMZ was created by firewalls from two different vendors, it is using defense diversity. Both firewalls are susceptible to vulnerabilities, and both are regularly updated as vulnerabilities become known.

At any time, an attacker could discover a previously unknown vulnerability on either firewall. If both firewalls are the same model from the same vendor, both firewalls will be susceptible to the same vulnerability at the same time. However, if both firewalls are

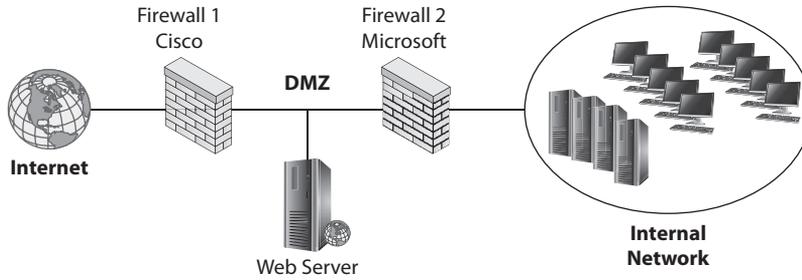


Figure 4-8 Implementing defense diversity with a DMZ

from different vendors, it is less likely that both will be subject to a particular vulnerability at the same time. Defense diversity with the DMZ firewalls provides an extra layer of protection for the internal network.



EXAM TIP A DMZ is typically created with two firewalls. Using firewalls from two separate vendors provides defense diversity for the DMZ.

Additionally, defense diversity requires the attacker to have more skills and knowledge in order to exploit both firewalls. For example, an attacker may have extensive expertise with the Cisco router, but very little on the Microsoft router. Fewer attackers will have the same level of expertise on firewalls from different vendors.

Comparing Network-based and Host-based Firewalls

As previously mentioned, firewalls are often classified as either network-based or host-based. A network-based firewall provides protection for a network, while a host-based firewall provides protection for a single host. As an example, consider Figure 4-9. The two firewalls used to create the DMZ are network-based firewalls and are filtering all traffic coming into or out of the network.

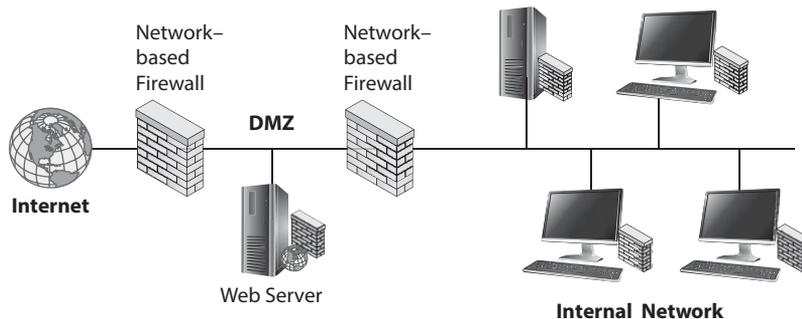


Figure 4-9 Network with host-based and network-based firewalls

Additionally, each internal system is running a host-based firewall that filters traffic coming into or out of the computer's network interface card (NIC). Both servers and desktop operating systems can run a host-based firewall as an additional application. As an example, many antivirus vendors now include firewalls as part of a comprehensive security suite.



TIP Some operating systems include built-in host-based firewalls running as additional software. For example, current Microsoft desktop and server products include the Windows Firewall, which is enabled by default.

Most organizations use both network-based and host-based firewalls. The network-based firewall filters and blocks malicious traffic coming from the Internet. Internal host-based firewalls block malicious traffic that reaches the internal network.

For example, imagine a user's home computer is infected with malware. When transferring files from home to work with a USB flash drive, the user can inadvertently infect the work computer just by inserting the USB flash drive into the work computer. This virus may then have a worm component that tries to locate and infect other systems on the network. However, if other systems have host-based firewalls installed, they can block the worm.

Exploring Remote Access Solutions

Remote access solutions allow users from remote locations to access internal network resources. Remote users can connect via a dial-up connection or a virtual private network (VPN), as shown in Figure 4-10. A VPN provides access to a private network over a public network such as the Internet.

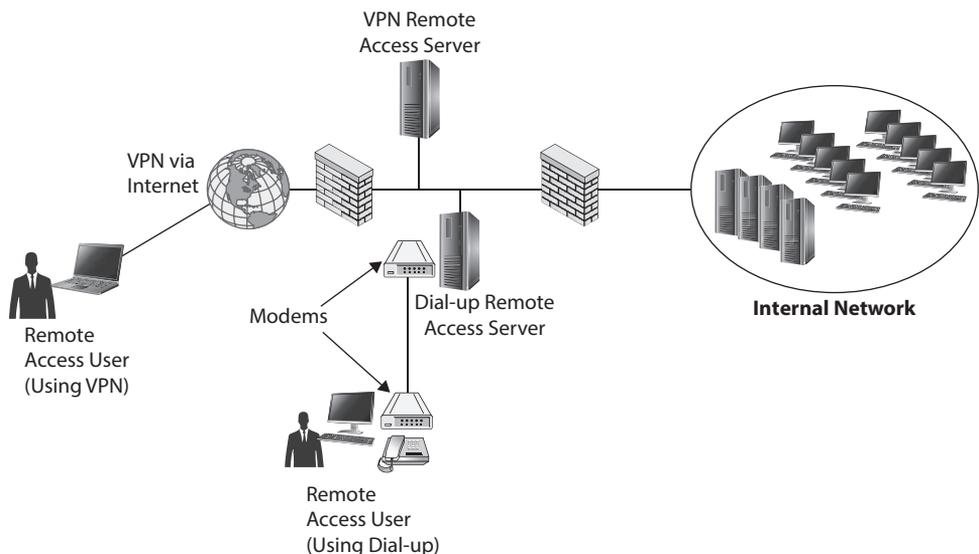


Figure 4-10 Remote access via dial-up or a VPN

In a dial-up connection, the user needs a modem and access to a phone line, and the remote access server also has a modem hooked up to a phone line. The user dials in, authenticates, and then accesses the private internal network. The telephone line is slow, with a PSTN line maxing out at less than 56 Kbps. Even an ISDN typically only provides 128 Kbps. In contrast, a broadband Internet connection, such as through a cable modem, provides much higher speeds. A quick speed test on my broadband connection shows a download speed of 37.18 Mbps and an upload speed of 6.78 Mbps.

Users with an Internet connection can connect via a VPN. The remote access client has a public IP address connected to the Internet, and the remote access server also has a public IP address assigned. The VPN client connects to the VPN server over the Internet, authenticates, and then accesses the private internal network.

Risks and Vulnerabilities

One of the primary risks of a remote access server is that anyone with access to a phone or the Internet can try to connect to the server. However, if strong authentication mechanisms are used, attackers cannot access the internal network because they lack the credentials to authenticate. Several different authentication protocols are available and they are explored later in this chapter.

VPNs transfer data over the Internet in an encrypted format to prevent attackers from capturing the data. This includes authentication information, such as usernames and passwords, and the actual data. If the data isn't encrypted, attackers can use a sniffer to capture the data transmissions and read the data. VPN tunneling protocols use encryption to prevent unauthorized disclosure of credentials and data.

Due to the inherent risks of using VPNs, many organizations simply don't implement them. Instead, users perform all their work on the organization's network while they are physically at the organization's location.



NOTE Many organizations are recognizing mutual benefits when workers are able to work from home. For example, there are fewer costs for the worker for the commute, and in some instances, the company can reduce its workspace. A VPN provides a good way for workers to telecommute, but the VPN must be secure.

Tunneling Protocols

A tunneling protocol encapsulates data used on an internal network and transmits it over the public network. For example, an internal network may use the Server Message Block (SMB) protocol with TCP/IP to access files and printers within an internal network. However, SMB is not used on the Internet. The tunneling protocol encapsulates SMB traffic to transmit it over the Internet.



NOTE SMB is used with other protocols within the TCP/IP suite. The supporting TCP/IP protocol information used on the remote network is also encapsulated with SMB.

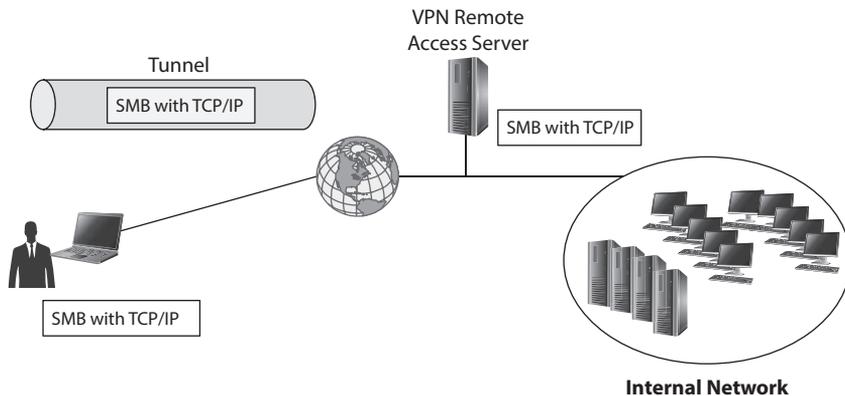


Figure 4-11 Encapsulating SMB data before sending it over the Internet

Figure 4-11 shows the overall process. The client wants to transfer data using SMB. Before sending the data via the VPN, the client software encapsulates the SMB protocol data within a tunneling protocol. This encapsulated SMB data is then transmitted over the Internet to the VPN server. The VPN server strips off the tunneling protocol data, leaving only the SMB and TCP/IP data. The SMB data is then sent to the appropriate server within the internal network using TCP/IP. When the data is returned to the client, the VPN server again encapsulates the data in a tunneling protocol to send it back to the VPN client.

Although Figure 4-11 may give you the impression that the SMB data is sent in plaintext, that is not the case. Instead, the tunneling protocol encrypts the data before sending it over the Internet. If attackers intercept any of the data, they will not be able to read it.

Some of the tunneling protocols that you may see in the SSCP exam to encapsulate VPN traffic are Secure Shell (SSH), Layer 2 Forwarding (L2F), Point-to-Point Tunneling Protocol (PPTP), Internet Protocol security (IPsec), Layer 2 Tunneling Protocol (L2TP), with or without IPsec, and Transport Layer Security (TLS). The following subsections briefly describe these protocols.

Secure Shell

The Secure Shell (SSH) protocol can be used to create an encrypted tunnel. Although the concept is similar to that of other tunneling protocols where the underlying protocol is encapsulated, SSH isn't commonly used for VPNs. It is used to encrypt other protocols such as Telnet and FTP.

Layer 2 Forwarding

Cisco created the Layer 2 Forwarding (L2F) protocol as an early VPN tunneling protocol. L2F provides an authentication mechanism, but does not encrypt the traffic. It is rarely used today.

Point-to-Point Tunneling Protocol

Point-to-Point Tunneling Protocol (PPTP) was much more popular than L2F. RFC 2637 defines the standard generically, and Microsoft developed a Microsoft version. The RFC version doesn't provide native encryption of the data, but the Microsoft version uses Microsoft Point-to-Point Encryption (MPPE) to encrypt the traffic in the tunnel. PPTP uses Generic Routing Encapsulation (GRE) to encapsulate the packets.

A significant weakness of PPTP is that the initial authentication connection is not encrypted. Information in the initial session can include usernames and passwords that can be intercepted by attackers. Even though the passwords are sent in a hashed form instead of in cleartext, attackers can capture them with a sniffer and use the usernames and password information in a replay attack. PPTP uses TCP port 1723.



TIP In a replay attack, the attacker captures a valid transmission, manipulates it in some way, and then resends it. If the attacker is able to capture reusable credential information, the attacker is able to impersonate the original sender but with different data.

Layer 2 Tunneling Protocol

Layer 2 Tunneling Protocol (L2TP) combines the strengths of L2F and PPTP and is documented in RFC 2661. Although L2TP doesn't encrypt the traffic itself, it is commonly combined with IPsec for encryption and authentication and referenced as L2TP/IPsec. L2TP/IPsec provides a high level of security for VPNs and has been widely implemented.

A primary drawback with L2TP/IPsec is that IPsec packets cannot pass through a system using NAT. Recall that NAT translates private IP addresses to public addresses and public IP addresses back to private. Because of how the translation is done, it corrupts the IPsec packets. L2TP uses UDP port 1701.



TIP NAT Traversal (NAT-T) was created to overcome the problem of NAT corrupting IPsec traffic. When equipment supports NAT-T, it's possible to send IPsec through a device using NAT-T.

Internet Protocol Security

Chapter 3 introduces Internet Protocol security (IPsec) as a secure protocol that provides mutual authentication of systems, integrity, and confidentiality. Recall that it uses the Authentication Header (AH) for authentication and integrity, and the Encapsulating Security Protocol (ESP) for encryption. It can be used with L2TP (as L2TP/IPsec) or by itself as a tunneling protocol. IPsec can work in two different modes:

- **Tunnel mode** In tunnel mode, the entire packet is encrypted and encapsulated into an IP packet. When tunnel mode is used for a VPN, this encapsulated packet is then sent over the Internet. When it reaches its destination, the encapsulated packet is decrypted and used on the internal network.

- **Transport mode** In transport mode, only the payload (the actual data) is encrypted. In other words, the IP header is not encrypted. In contrast, tunnel mode encrypts the entire packet, including the IP header. Transport mode is used within an internal network to encrypt traffic from one point to another.

The National Institute of Standards and Technology (NIST) published Special Publication (SP) 800-77, *Guide to IPsec VPNs*. It provides significantly more details on using IPsec for VPNs. You can view a list of SP 800 documents here: <http://csrc.nist.gov/publications/PubsSPs.html>.

Transport Layer Security

Many VPN solutions use TLS for the VPNs. As you might recall from Chapter 3, TLS is commonly used to encrypt HTTP (as HTTPS) over TCP port 443. A TLS VPN allows clients to connect with a web browser and also uses TCP port 443.



NOTE TLS VPNs originally started as Secure Sockets Layer (SSL) VPNs, and you'll often see them marketed as SSL VPNs even though they support both TLS and SSL. However, TLS is the designated replacement for SSL, and most organizations no longer use SSL due to known vulnerabilities.

A TLS VPN is a good solution when the VPN is having trouble getting through a NAT device, such as when using L2TP/IPsec without NAT-T. It's also becoming popular because it doesn't require any additional configuration on the client. The client only needs a web browser and knowledge of the URL to connect to the organization's TLS VPN.

NIST published SP 800-113, *Guide to SSL VPNs*, in 2008. It mentions that an SSL VPN works the same as a TLS VPN, but at the time it was written, "SSL VPN" was a more popular term. Still, SP 800-113 provides a good in-depth description of TLS VPNs.



EXAM TIP An SSL VPN uses port 443 and allows clients to access internal resources with a web browser.

Authentication

One of the primary methods of controlling who can access a remote access server is authentication. Users claim an identity (such as with a username) and prove their identity (such as with a password) during an authentication process. Based on a user's proven identity, access controls are checked to see whether the user is authorized to use the remote access server. If a user cannot be authenticated, the remote access server denies access. The following sections identify some common remote access authentication protocols.

PAP

The Password Authentication Protocol (PAP) is an older authentication protocol. Usernames and passwords are passed over the network in cleartext, which makes them highly susceptible to sniffing attacks. PAP is rarely used today.

CHAP

The Challenge Handshake Authentication Protocol (CHAP) is more secure than PAP because the actual password is never sent over the wire. Instead, a *nonce* (a number used once) is combined with a shared secret known only to the client and the remote access server. The result of combining the nonce and the shared secret is then hashed with Message Digest 5 (MD5).

Hashing was briefly introduced in Chapter 1. As you'll recall, a hashing algorithm (such as MD5) creates a number. More specifically, MD5 creates a 128-bit number.

Putting it all together, here's a big-picture view of how the process works. The client sends an authentication request to the remote access server. The server responds by sending back the nonce. The client then combines the shared secret with the nonce and hashes the result with MD5. This MD5 is then sent back to the server. Because the server knows the nonce (it created it) and it knows the shared secret, it can also calculate the hash on the server. If the MD5 hash sent by the client matches the MD5 hash calculated on the server, then the client is authenticated.

MS-CHAPv1 and MS-CHAPv2

Microsoft improved CHAP with its versions of Microsoft CHAP (MS-CHAPv1) and MS-CHAPv2. Although MS-CHAPv1 is no longer used, MS-CHAPv2 is still used in many Microsoft implementations.

A primary improvement of MS-CHAPv2 is mutual authentication. In other words, the client authenticates to the server and the server also authenticates back to the client. This helps prevent the client from sending authentication data to an attacker impersonating the remote access server.

EAP

The Extensible Authentication Protocol (EAP) provides a more secure authentication process than many other remote authentication methods. EAP doesn't use a specific method, but instead security developers can extend existing authentication methods using EAP. The following list identifies some of the commonly used EAP methods:

- **Protected EAP (PEAP)** PEAP encapsulates and encrypts the EAP transmission in a TLS tunnel. PEAP requires a certificate on the server, but clients are not required to have a certificate. A major advantage of PEAP is support by Microsoft systems and MS-CHAPv2 often uses PEAP.



TIP Chapter 14 covers certificates in more depth, but as an introduction, certificates support strong authentication and encryption services. A certificate authority (CA) issues certificates, and systems can query the CA to verify the validity of a certificate.

- **EAP-Tunneled TLS (EAP-TTLS)** EAP-TTLS also uses TLS to create a tunnel to encapsulate the authentication method. It exchanges authentication information as attribute-value pairs (AVPs), allowing it to secure many older authentication

methods (such as PAP, CHAP, and MS-CHAP). Note that while PAP sends the authentication information in cleartext, the TLS tunnel ensures that the authentication information is encrypted while in transit. EAP-TTLS requires a certificate on the authentication server, but clients do not need a certificate.

- **EAP-TLS** EAP-TLS provides the best security of these methods. It requires the server and each client to have a certificate, providing mutual authentication.

RADIUS

The Remote Authentication Dial-in User Service (RADIUS) provides centralized authentication, authorization, and accounting (AAA) for remote clients. When a remote user attempts to connect, the remote access server passes the request and user credentials to a RADIUS server. The RADIUS server verifies the user's credentials and verifies the user is authorized to connect to the remote access server. After the user connects, the RADIUS server tracks the user activity in an accounting log.

Even though it has “dial-in” in the name, RADIUS is used for remote clients that connect through other sources, such as through a VPN.



EXAM TIP RADIUS is not only for dial-up. It also provides AAA services for VPN remote access connections. Other implementations include configuring RADIUS as an 802.1x authentication server to require authentication for wireless clients, commonly known as WPA2-Enterprise.

Although a RADIUS server can be used for a single VPN server, it is more commonly used with multiple remote access servers. For example, Figure 4-12 shows several VPN servers hosted by an organization in different cities. Remote access users can connect to any of the servers over the Internet and then have access to the network physically located in that city.

In the figure, a remote access user is connecting to the Virginia Beach VPN server over the Internet. The user passes authentication credentials to the VPN server, and the VPN server then passes this authentication information to the RADIUS server.



TIP When a VPN server is using a RADIUS server for authentication, it is called a RADIUS client. In this context, the VPN server acts as the server for the VPN client but as a client for the RADIUS server.

The RADIUS server might have an authentication database with the credential information (such as a database of usernames and passwords) of all authorized users. It's also possible for the RADIUS server to check authentication with another server. For example, in a Microsoft network, the RADIUS server passes the authentication information to the domain controller. If the credentials are verified, the information is passed back to the remote access server and the client is authenticated. If a user connects to one of the other VPN servers (such as in New York), the VPN server passes the authentication information to the RADIUS server and the rest of the process is the same.

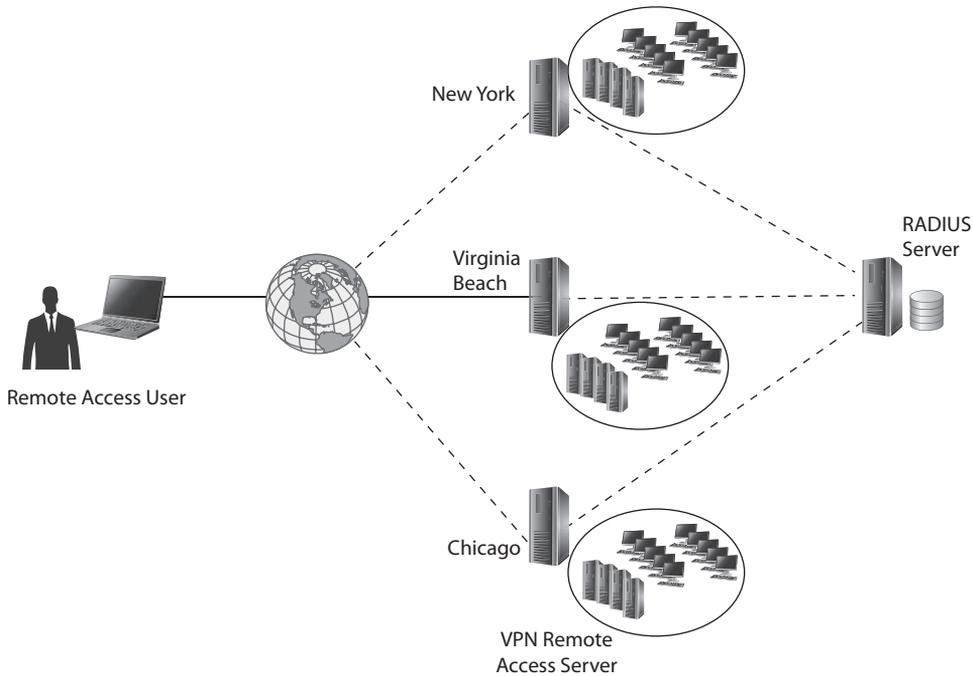


Figure 4-12 Using RADIUS for authentication

A primary benefit is that the authentication is centralized and only a single authentication database needs to be maintained. Imagine for a moment that an organization uses VPNs to allow administrators to remote in and perform administrative work. Administrators can remote into any of ten different VPN servers located in offices around the country. When a new administrator is added or removed from the administrative team, only one entry is needed on the RADIUS server. In contrast, if RADIUS were not used for central authentication, each of the ten different VPN servers would have to be modified.

Historically, RADIUS has used UDP and it only encrypts the password during the authentication session. However, RFC 6613 provides specifications for using TCP, and RFC 6614 provides specifications for using TLS to encrypt the entire authentication session. As I write this, both of these RFCs are still in the experimental stage.

Diameter

Diameter is another AAA protocol and it was created as an alternative to RADIUS. It isn't backward compatible with RADIUS, but it does provide several improvements. It uses TCP instead of UDP and supports both IPsec and TLS to encrypt sessions.

The name Diameter implies that it is twice as good as RADIUS. The radius of a circle is the distance from the center to an edge of the circle. The diameter of a circle is the distance from one edge of the circle to the other edge of the circle. In other words, the diameter of a circle is twice the distance of the radius.

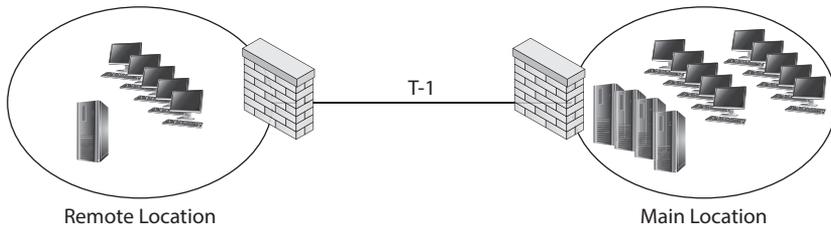


Figure 4-13 Site-to-site VPN

TACACS+

Terminal Access Controller Access Control System+ (TACACS+) is another AAA protocol and is an alternative to RADIUS. Cisco originally created it as a proprietary protocol but later released it as an open standard. TACACS+ uses TCP so communication is more reliable than RADIUS transmissions. Also, TACACS+ encrypts the entire authentication process, while RADIUS only encrypts the password. TACACS+ uses TCP port 49.



EXAM TIP TACACS+ encrypts the entire authentication session. In contrast, RADIUS encrypts only the user's password.

Traffic Shaping

Traffic shaping refers to delaying certain types of traffic so that other traffic performs better. For example, audio and video streaming can consume a significant amount of bandwidth, so a traffic shaping policy would delay this type of traffic, giving preference to other traffic. Note that traffic shaping is not the same as packet filtering. Packet filtering drops undesirable traffic. Traffic shaping only delays it. The result is that other traffic is given priority.

Wide area networks (WANs) often use VPNs to connect, and traffic shaping is useful to optimize the WAN traffic. As an example, imagine that a remote office connects to the main location for a business via a site-to-site VPN as shown in Figure 4-13. Instead of connecting over the Internet, the locations are connecting via a leased T-1 at a speed of 1.544 Mbps. Compared to internal speeds of 1 Gbps or more, 1.544 Mbps is quite slow. Worse, the performance of this link can become even slower if users are regularly using it for streaming audio and video.

Traffic shaping will optimize the WAN link by delaying undesirable traffic. Note that this example focused on streaming audio and video, but an organization can choose to delay any type of traffic based on its needs.

Access and Admission Control

Network access control (NAC) is a technology used to control which clients are granted access to a network. This is especially important when remote access solutions allow external clients into the private network. However, NAC can also control unhealthy clients within the network.

The overall goal of NAC is to inspect clients regularly to ensure they meet specific *health* requirements. Health is defined by criteria such as having up-to-date antivirus software installed and running, having a host-based firewall enabled, and keeping the operating system up to date. These elements are defined in a health policy specified by administrators.

When NAC identifies an unhealthy client, it typically quarantines the client to prevent it from infecting other computers on the network. A quarantined network would also include resources that clients can use to improve their health, such as updates for the antivirus software or operating system.

When computers are completely controlled within an organization, administrators have several different tools and methods they can use to ensure that the systems are secure. They can automate the deployment of antivirus software and updates to the clients, automate the deployment of operating system updates, control what software is installed and running on the systems, and even control what the user can do on the system, such as what software the user can install.

However, remote access users may use a home computer over a VPN connection to access the private network. Because it's a home computer, the organization's administrators don't have direct control over the computer. A significant risk is related to malware. If the infected system connects to the internal network, it can easily infect other computers on the network.

Figure 4-14 shows the overall process of how a NAC solution works. After a remote access user connects (via either dial-up or VPN connection), the remote access server

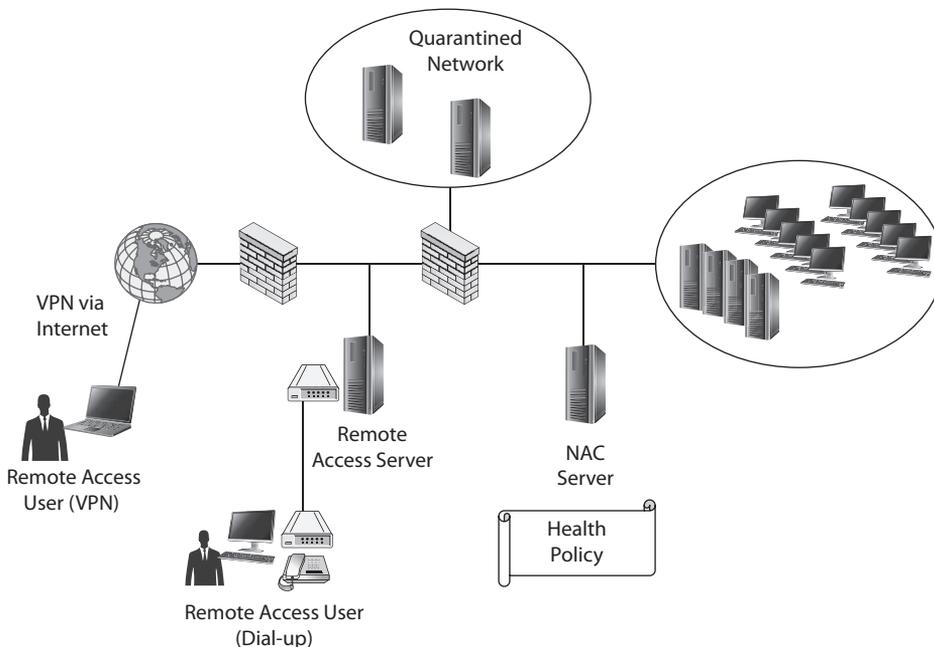


Figure 4-14 Using network access control

coordinates with the NAC server to determine the health policy for the organization. The client is inspected to determine whether it meets the criteria of the health policy. If the client is healthy, it is granted access to the network. If the client is not healthy, the NAC server restricts access for the client to the quarantined network only.

Although NAC is often used for remote clients, it can also be used for internal clients. This helps prevent unhealthy systems from staying operational on the internal network. As an example, many automated methods exist to deploy operating system updates to clients. However, there are multiple reasons that a critical update may not be installed on a system, such as if the system is powered off while a user is on an extended vacation, business trip, or leave of absence. When the user turns the computer back on and tries to access the network, NAC checks its health and quarantines it if it is not healthy.

Exploring Virtual Environments

Virtualization refers to replacing hardware with software. One of the first ways organizations did this was with servers. For example, imagine an organization has 100 servers within a data center. The organization can replace these servers with 20 powerful servers. Each of these physical servers can host five virtual machines (VMs) so the organization still hosts 100 servers.

While the cost of the 20 servers is high, an organization will quickly recoup its costs through other savings. For example, these 20 servers require less power; less heating, ventilation, and air conditioning (HVAC); and less space. Additionally, the organization reduces the hardware costs in future hardware refresh cycles. Instead of replacing 100 computers every three to five years, the organization has to replace only 20 servers.



NOTE The Gartner group reported in 2014 that at least 70 percent of x86-based servers were virtualized. VMware and Microsoft own the majority of the market share in virtualization. Gartner also reported that more than 90 percent of customers reported using VMware as their primary virtualization software, while 48 percent reported using Microsoft's Hyper-V as their secondary virtualization software.

In addition to creating VMs, it's possible to create virtual networks within a host machine. These aren't the same as VLANs created on switches. Instead, the host machine allows you to create and configure virtual switches and virtual routers running on the host machine. Administrators configure these virtual network devices to connect the host machines together and connect them to the live network.

Virtualization Terminology

Three important terms related to virtualization are host, guest, and hypervisor. The physical server is the host machine. Each VM running on the host machine is a guest, and a host can run multiple guest VMs. The host runs hypervisor software such as Microsoft Hyper-V or VMware vSphere. The hypervisor is the software on the physical server that

creates, runs, and manages the VMs. It also manages the resources for each of the VMs. These resources include the virtual data storage, memory, and networking components.

Because a host runs multiple guest VMs, it requires more resources than a typical server does. These physical resources include multiple processors, a significant amount of RAM, and fast drives. The host system shares the use of these resources with all of the VMs.

Shared Storage

Guest VMs on host machines often share the same data storage, so it's important to ensure it has enough space and is fast. At a minimum, it should use a fast Redundant Array of Independent Disks (RAID) subsystem. Another option is to use solid-state drives (SSDs), which are much faster than traditional hard disk drives.

A security concern with using shared storage is confidentiality. You want to ensure that users within one VM are not able to access data in another VM. Strong access controls will typically provide enough protection to prevent unauthorized access.

Chapter 11 covers data classifications in more depth. For example, the U.S. government classifies data as Top Secret, Secret, Confidential, and Unclassified. With this in mind, it's best to put guest machines with similar data classifications on the same host system. In other words, you wouldn't put a guest machine hosting Top Secret data on the same system hosting only Unclassified data.

Virtual Appliances

A virtual appliance is a preconfigured VM with an operating system and an application. The VM is ready to be added to a host machine as a guest. The benefit of a virtual appliance is that it is just a set of files that administrators can copy or import to a host machine. Of course, these virtual appliances are created for specific hypervisors. For example, a VM created for a Microsoft Hyper-V hypervisor will not run on a host machine using a VMware vSphere hypervisor.



NOTE Typically, *appliance* refers to a hardware device used for specialized tasks. For example, a UTM appliance is a hardware device that includes several security features on a single device. In contrast, a virtual appliance is software running as a guest VM on a host server.

As an example, Openfiler has created virtual appliances for file-based Network Attached Storage and block-based Storage Area Networking (NAS/SAN). The NAS features include support for the Common Internet File System (CIFS) and Network File System (NFS). The SAN features include support for Internet Small Computer System Interface (iSCSI) and Fibre Channel (FC). Openfiler has created virtual appliance images for a free version, and a full-featured Commercial Edition.

While Openfiler is one example, there is almost no limit to what vendors can create as a virtual appliance. This has caused many organizations to look at servers and virtualization differently. Instead of evaluating how many servers they need, they evaluate how many applications they want to host as VMs.

Continuity and Resilience

Virtualization also contributes to continuity and resilience, increasing overall availability. Many cloud providers include elasticity features with virtual servers that they rent. For example, Amazon Elastic Compute Cloud (Amazon EC2) makes heavy use of virtualization. When customers rent cloud services through Amazon EC2, they rent access to virtual systems.

One of the benefits of elasticity is that capacity can easily shrink or expand based on the demand. In this way, users pay only for what they're using. For example, if a customer's website suddenly has a burst of activity and needs additional processing power or memory, it's relatively easy for EC2 to expand the available resources for the site. When the additional resources are no longer needed, the virtual system can shrink back to the original size.

Many hypervisors allow administrators to manage the guest machines using a point-and-click interface. This allows them to easily create backups of an entire VM with just a few clicks. Additionally, some hypervisor management tools allow administrators to move VMs between host systems via a central server. As an example, imagine one host system is running five VMs but increased usage is overloading the host system. Administrators can use a point-and-click interface to move one or more of the VMs to another host, freeing up resources on the original host system.

It's also much easier to restore a VM if it crashes, when compared to a physical server, assuming administrators keep up-to-date backups of the VMs. If a VM crashes, administrators can restore it by restoring the files. In contrast, if a physical server crashes, it takes much more to rebuild it from scratch.

Similarly, most hypervisors allows you to create snapshots of a system. These are very useful, especially before performing a risky task such as an application upgrade. If the upgrade causes problems, you can revert the system to the state it was in when you created the snapshot.

Separation of Data Plane and Control Plane

Traditional hardware routers implement both the data plane and the control plane. A current trend is to replace hardware routers with software-defined networking (SDN) to separate these two planes. Before discussing SDN concepts, it's worth defining the data plane and the control plane:

- **Data plane** The data plane (sometimes called the forwarding plane) uses rules within an ACL to identify whether the router will pass traffic or block traffic. These are the same type of rules described in the "Packet-Filtering Firewall" section earlier in this chapter. The data plane is responsible for deciding if traffic should be forwarded or not. The data plane is proprietary. In other words, rules used by one router often have different syntax when compared to rules created by another router manufacturer.
- **Control plane** The control plane is responsible for identifying paths to other networks. Routers typically do this with routing protocols such as OSPF and BGP. These routing protocols allow routers to share information with each other to create a virtual map of the network.

Software-defined Networking

Software-defined networks separate the control and data planes using virtualization technologies. This allows an organization to replace many of the hardware routers and switches with software. This helps organizations move away from proprietary hardware and also gives them much more flexibility.

As an example, Chapter 2 discusses Attribute-based Access Control (ABAC) models, which SDNs typically support. An ABAC model allows administrators to create policies that are much more versatile than rules within an ACL.

While software-defined networking (SDN) and network function virtualization (NFV) are relatively new technologies, they are growing fast. The 2015 Edition of the SDxCentral *SDN and NFV Market Size Report* includes several predictions of explosive growth. For example, it indicates sales of SDN products replacing traditional layer 2 and layer 3 hardware (switches and routers) will exceed \$10 billion in 2015 and grow to more than \$105 billion in 2020. For comparison, SDxCentral reported the SDN market size in 2013 was about \$1 billion.

Attacks and Countermeasures

As additional software running on the server, the hypervisor represents another potential risk. If an attacker can access the host machine from within a VM, the attacker potentially can access all of the guest machines. The following list identifies some of the specific risks associated with virtualization:

- **Improperly configured hypervisor** The security of all the guest systems is directly related to the security of the hypervisor on the host. If administrators don't adequately restrict access, it becomes easier for unauthorized personnel to gain access to all the systems. The key here is to ensure that administrators have appropriate training to manage and secure the hypervisors.
- **Infected virtual appliances** Free virtual appliances sound very appealing. However, it's possible for attackers to infect these with Trojans or other malware, and once you add the virtual appliance, attackers have a foothold in your environment. This is similar to the many new computers built in foreign countries that have been shipped to users with malware preinstalled. Just as a single infected physical machine in your network can spread malware to other systems, a single infected VM can also spread malware. A key here is to be suspicious of any files downloaded from the Internet and to only use trusted files from a trusted source.
- **Data leakage** Because VMs are simply files, it's very easy to copy an entire VM just by exporting and copying the files. If unauthorized personnel have access to these files, they might be able to import them to their own system running a hypervisor and access all of the data within the VM. Additionally, attackers might be able to scan the data within the files and extract meaningful data. These VM files should be protected with the same level of security as the highest classification of data stored on the VM.

- **Improperly configured virtual network** Most network devices are configured by administrators with specialized knowledge about routers and switches. However, server administrators typically configure virtual network devices (such as virtual routers and virtual switches) on the host. They might not have the same level of knowledge and expertise and might configure them improperly. An improperly configured virtual network might incorrectly direct sensitive traffic to guest VMs on the host system.
- **VM escape** A VM escape is an exploit that allows a guest to interact with the host system. As an example, security specialists identified a vulnerability in VMware that allowed a VM escape in 2008. The security company Immunity implemented code into its CANVAS penetration toolkit to exploit this vulnerability in 2009. Patches released by VMware eliminate that specific VM escape vulnerability, and patching is a key countermeasure for any VM escape vulnerability. When they're discovered, the vendor releases patches, so it's best to keep the host server and guest systems up to date with current patches.



EXAM TIP VM escape is a known attack against virtual systems. If successful, an attacker can access the host system and all virtual systems within the host.

Understanding Cloud Computing

Cloud computing refers to any type of computing services provided over the Internet. It has become quite popular in recent years. Cloud computing is very cheap and sometimes even free, but it includes some security concerns.

Three common services provided through the cloud are Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS). The lines between these cloud services are sometimes blurred, partly due to how cloud vendors market them. However, one way to categorize them is to look at the different maintenance responsibilities, which correlate with security responsibilities. The *Cloud Computing Security Requirements Guide*, published by the U.S. Department of Defense (DoD), includes a figure similar to Figure 4-15. Figure 4-15 shows that the vendor has the most maintenance and security responsibility for SaaS, a little less for PaaS, and the least for IaaS.



EXAM TIP Cloud computing includes Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS), and it's important to know the difference between these three services. Customers have the most maintenance and security responsibility for IaaS and the least responsibility for SaaS.

Software-as-a-Service (SaaS) is also known as on-demand software. It provides users with access to software or applications over the Internet. For example, Google provides

Figure 4-15
Maintenance
and security
responsibility

SaaS	Vendor	Customer
PaaS	Vendor	Customer
IaaS	Vendor	Customer

several applications through the cloud, including Gmail for e-mail and Google Docs for other types of files, such as spreadsheets, presentations, and word processing documents. Users access the applications using a web browser. This allows the cloud-based applications to support different operating systems.

Domain Name System (DNS) also fits into the SaaS service model. Chapter 3 discusses DNS and how it resolves host names to IP addresses. For example, if you want to reach Google, you type `google.com` as the address in your web browser, and your system queries DNS to return an IP address. Organizations that maintain resources on the Internet need to ensure that they register and maintain records for these resources. They could create and maintain their own DNS server, but it's often easier to outsource this and use a cloud provider to manage the records on the cloud provider's DNS server. A huge benefit is that the cloud provider manages the security of the DNS server and often has more knowledge and expertise to thwart DNS attacks.

The “Exploring Remote Access Solutions” section earlier in this chapter discussed the use of VPNs for remote access. However, VPNs can also be used in a SaaS model. Users connect to the VPN service provider and surf the Internet through it. The VPN encrypts all traffic and prevents attackers from capturing the traffic. This is especially useful for wireless users accessing Internet resources through public hotspots.

Similarly, a proxy server is typically used within an organization for caching and web filtering. However, you can also subscribe to proxy services from some cloud providers. Administrators configure the proxy settings within the web browsers, which causes all Internet access to go through the cloud-based proxy server. The proxy server performs web filtering just as an in-house proxy server does.

Platform-as-a-Service (PaaS) provides users with a computing platform. This includes hardware, an operation system, and applications. As an example, many web-hosting companies rent access to servers that are preconfigured with an operating system and web hosting applications, such as WordPress. Customers may install the applications from a list of available options, or the vendor may install the applications. The vendor performs all of the hardware maintenance on the server, and typically performs the majority of the software maintenance. For example, the vendor would keep the server up to date with current patches. The customer is responsible for keeping applications up to date, but the vendor may provide assistance.

Infrastructure-as-a-Service (IaaS) is sometimes called Hardware-as-a-Service, and this is a good way to think of it. Customers rent access to hardware, such as servers and networking infrastructure, and the vendor maintains the hardware. The customer is responsible for installing operating systems and applications on the hardware, and is responsible for maintenance of the operating systems and applications.

Cloud Operation Models

When considering security with cloud-based computing, it's also worth examining the different operation models, or deployment models. NIST SP 80-144, *Guidelines on Security and Privacy in Public Cloud Computing*, defines the following four models:

- **Public cloud** Public cloud-based services are provided by third-party vendors and are available to anyone. As an example, Apple provides cloud-based storage via its iCloud service.

- **Private cloud** A private cloud is available only to users within an organization. As an example, a company can set up cloud-based storage for its employees. The company would ensure the server is accessible via the Internet, but employees would be the only personnel authorized to access the server.
- **Community cloud** A community cloud is a private cloud that is shared by two or more organizations. For example, imagine that Globex set up a private cloud for Globex employees. Later, Globex and Acme sign a partnership deal and need an easy method to share files. Globex can configure its private cloud so that Acme employees can access it. It is now a community cloud.
- **Hybrid cloud** A hybrid cloud is a combination of any two or more clouds. It can be a combination of different clouds or similar clouds. For example, imagine that both Globex and Acme host their own private clouds. Later, they sign a partnership deal with each other, and Acme shares part of its private cloud with Globex, creating a community cloud. The Globex employees will be able to access a single cloud-based service and have access to both the private Globex cloud and the Acme community cloud.

A private cloud provides the best security for sensitive data. However, it's worth remembering that a private cloud is still accessible via the Internet. If users don't use secure passwords, or if they reuse their credentials on other sites, attackers may be able to access the data in the cloud. While a community cloud is similar to a private cloud in that it is private between the companies sharing the data, it's important to ensure that only shared data is stored in the community cloud.

Storage

Storage in the cloud is almost limitless. An organization can easily rent the storage space it needs, and that space is often very cheap or even free. For example, Google Drive provides every user with 15GB of storage space for free. If you need more, you can rent 100GB for \$1.99 a month. Microsoft offers unlimited OneDrive storage for Office 365 subscribers. This is a moving target though, and might be different by the time you read this.

However, free storage isn't necessarily safe storage. Organizations need to be careful about what they store in the cloud. A data breach with the cloud provider can result in a data breach for an organization.

When picking a cloud storage location, it's sometimes useful to consider the legal jurisdiction of the storage location. As an extreme example, imagine the cloud provider accesses an organization's data and sells it to a competitor. If the cloud provider is located in the same country, the organization can pursue a legal remedy. However, if the organization is in the United States and the cloud provider is in Russia, the organization may not have any legal options. Similarly, if a U.S. cloud provider outsources storage locations to other countries, sensitive data might wind up being stored in unfriendly locations.

Electronic discovery (e-discovery) refers to any process that searches data with the goal of using it as evidence. It also includes locating and securing the data so that it can be used in a civil or criminal case. Using a cloud provider from a local jurisdiction ensures that e-discovery processes can be used when necessary.

Privacy

Privacy of data posted on the Internet has always been a concern. A good guideline to adopt is that if you post data on the Internet, you have lost control of that data and might have also lost privacy of that data. This includes data that you store in the cloud.

Cloud computing attempts to provide users with privacy of their own data, but any cloud provider is vulnerable to attacks, potentially resulting in data loss or data leakage. When an organization hosts data on the Internet through a cloud provider, protection of that data is now partially dependent on the provider using adequate protection. If the provider has a failure, it can result in a data breach.

For example, in June 2011, Dropbox, a cloud-based storage provider, did an update that accidentally allowed anyone to log on to any account with any password. Users could have entered the wrong username or any password and suddenly had access to someone else's data. After the media reported about the mistake, Dropbox let users know about the problem and said that it only affected 1 percent of users. Although 1 percent sounds like a small amount, that's 1 percent of 25 million users, or 250,000 users.

Cloud customers also have a responsibility for security, and poor security practices by customers can lead to data breaches. In 2014, attackers accessed data on Dropbox again, but this was due to data breaches from other organizations and weak security practices of customers. Specifically, customers were using the same username and password on Dropbox that they used with other companies. When attackers gained access to user credentials, they methodically tried them on different services such as Dropbox.

Encryption is an excellent method of providing confidentiality. If you are using encryption to protect your data in the cloud, the data should be encrypted on the client end and not in the cloud. In other words, if you want to store sensitive data in the cloud, encrypt it before posting it to the cloud. If you use encryption tools made available by the cloud provider, it's very possible that the provider holds the decryption keys and employees of the provider can access the data or may inadvertently decrypt it, making it available to other users.

Data Control and Third-party Outsourcing

Organizations that store data in the cloud are effectively outsourcing some of the security requirements. With this in mind, an organization should consider implementing some security controls to increase the control of the data.

A service level agreement (SLA) is an agreement between an organization and a vendor. Organizations use SLAs when contracting services from service providers, including cloud service providers. An SLA typically stipulates performance expectations, such as minimum uptime, and can include expectations for the availability and security of data. The SLA can mandate certain requirements such as following:

- **Data portability** Data portability allows users to reuse data across interoperable applications. Cloud providers can implement different application programming interfaces (APIs) to support data portability.
- **Data destruction** When the data is no longer stored with the cloud provider, it's important that it is destroyed. Chapter 11 discusses various methods of

adequately destroying data, and it's important to ensure that the cloud provider uses similar methods to eliminate any data remnants of sensitive data.

- **Auditing** An organization owns the data, and this ownership doesn't change when the organization chooses to store the data with a cloud provider. Auditing processes can help ensure that employees of the cloud provider are not accessing the data inappropriately.
- **Data resilience** The cloud provider needs to implement appropriate security controls to ensure the resilience of the data. This includes appropriate backup and recovery processes to overcome any potential system failures.

Compliance

Organizations that must comply with laws and regulations need to be careful when using the cloud, especially for data storage. The security provided via the cloud may not be adequate to protect some data, but this doesn't relieve the organization of its legal responsibilities. For example, the U.S. Health Insurance Portability and Accountability Act (HIPAA) of 1996 mandates the protection of health-related data. This includes any data that directly relates to patients of medical facilities and any data collected about an individual's health. Many non-health-related organizations must comply with HIPAA because they collect information on users for insurance and health plans.

Organizations that process transactions from major credit cards need to comply with the Payment Card Industry Data Security Standard (PCI DSS). E-commerce transactions traverse the Internet, and data is often stored in cloud-based storage. Organizations that use cloud-based storage for credit card data need to ensure that the cloud-based provider meets the minimum standards defined by PCI DSS.

Some cloud providers are meticulous about complying with laws and standards, but other providers might not be so careful. Organizations need to be aware that they are ultimately responsible for protecting their data. If attackers gain access to sensitive data that an organization stores in cloud-based storage, the organization is still responsible for the data breach.

Chapter Review

Switches connect devices together in a network, and routers connect networks together within a local area network (LAN). Virtual local area networks (VLANs) increase security by segmenting traffic with a switch similar to how a router segments traffic between networks. Physical security methods protect routers and switches.

Telecommunications is the transmission of any type of signals for communicating. People communicate over the Internet via Internet service providers (ISPs) using methods such as the public switched telephone network (PSTN), integrated services digital network (ISDN), Digital Subscriber Line (DSL), cable, satellite, and wireless. Voice over Internet Protocol (VoIP) is used to transfer multimedia and voice communications over IP networks, including the Internet. Secure Real-time Transport Protocol (SRTP) provides confidentiality, authentication, and replay protection for VoIP signals.

Many organizations use their own private phone systems, or PBXs, which should be protected. It's common to protect the systems with physical security by locking them in secure server rooms. Logical protections, such as protecting the administrator password, restricting call forwarding, and restricting long distance calling, provide another layer of protection for phone systems.

Retrieved web pages are cached on the proxy server, and the proxy server serves these cached web pages instead of using Internet bandwidth to retrieve the web page from the Internet again. Website filtering allows a proxy to block a user's access to specific website locations. Most proxy servers also have NAT installed to translate private and public IP addresses.

Firewalls are an important security element in a network and for individual computers. Packet-filtering firewalls can filter traffic based on IP addresses, ports, and some protocols by using the protocol ID. Stateful inspection firewalls identify active TCP and UDP sessions and can open and close firewall ports dynamically based on the needs of active connections. An application firewall includes different elements to examine specific commands used by different protocols, such as HyperText Transfer Protocol (HTTP), File Transfer Protocol (FTP), and Simple Mail Transport Protocol (SMTP). Next-generation firewalls and unified threat management (UTM) devices add more filtering capabilities such as blocking malware, filtering spam, and URL filtering.

A DMZ is commonly created with two firewalls from separate vendors to achieve defense diversity. This is in line with an overall defense in depth strategy. Even if a vulnerability appears in one firewall, it's unlikely a vulnerability will appear in both firewalls at the same time. Additionally, to succeed in penetrating an internal network that employs defense diversity, an attacker must have in-depth knowledge about firewalls from two vendors.

Firewalls can be network-based or host-based. A network-based firewall protects traffic going into or out of an overall network, while a host-based firewall protects traffic for an individual system. Network-based firewalls are typically hardware-based, while host-based firewalls typically run as an additional software component on a server or desktop system. Using both network-based and host-based systems is part of an overall defense in depth strategy.

Many organizations provide access to their internal network through a dial-up or VPN remote access solution. A VPN provides access to the private network over a public network such as the Internet. VPNs use a tunneling protocol to encapsulate the protocols needed on the internal network. Tunneling protocols include Secure Shell (SSH), Layer 2 Forwarding (L2F), Point-to-Point Tunneling Protocol (PPTP), Layer 2 Tunneling Protocol (L2TP), Internet Protocol security (IPsec), and Transport Layer Security (TLS). When using TLS, users often only need a web browser to connect to the VPN server.

Authentication protocols used for remote access include Password Authentication Protocol (PAP), Challenge Handshake Authentication Protocol (CHAP), Microsoft CHAP (MS-CHAPv1), MS-CHAPv2, Extensible Authentication Protocol (EAP), Remote Authentication Dial-in User Service (RADIUS), Diameter, and Terminal Access Controller Access Control System+ (TACACS+). RADIUS, Diameter, and TACACS+ also provide authentication, authorization, and accounting.

Network access control (NAC) helps protect internal networks by differentiating between healthy and unhealthy systems and restricting access of unhealthy systems to quarantined networks. NAC is very useful with remote access because remote clients

aren't directly controlled by the organization and can have varying levels of security applied. NAC can also be used within the organization to ensure that unhealthy clients are isolated from healthy clients.

Virtualization refers to replacing hardware with software. Host systems include a hypervisor that creates, runs, and manages virtual machines (VMs), also known as guest systems. An important security step with virtualization is to keep all hosts and guests up to date with current patches.

Cloud computing includes Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS). Although cloud computing provides reduced costs for many users, it also includes increased risk. Data hosted in the cloud can be compromised due to errors or problems with the cloud provider. Encrypting data on the client before storing it in the cloud can help provide confidentiality of data, but encrypting data using a cloud provider's encryption services often does not provide adequate protection.

Questions

1. Which of the following best describes the mapping of data held within a switch's table?
 - A. IP address to port
 - B. MAC address to port
 - C. IP address to MAC address
 - D. Physical port to logical port
2. Which of the following is the best choice to segment traffic on a network?
 - A. VLAN
 - B. EAP
 - C. SSL
 - D. TLS
3. Which of the following can provide security for VoIP?
 - A. RADIUS
 - B. TACACS+
 - C. PSTN
 - D. SRTP
4. Your organization has a private phone system. Of the following, what is the best choice to control call forwarding?
 - A. Ensure that the administrator password is kept private and changed often.
 - B. Restrict phone numbers that can be used with call forwarding.
 - C. Restrict long distance calling.
 - D. Protect the phone system with physical security.

5. Of the following choices, what represents the primary benefits provided by a proxy server?
 - A. Caching and filtering
 - B. Authentication and caching
 - C. Authentication, authorization, and accounting
 - D. Stateful inspection
6. A packet-filtering firewall can block ICMP traffic, such as ping requests. How does a packet-filtering firewall identify ICMP traffic?
 - A. Based on the protocol ID having a value of 1
 - B. Based on the protocol ID having a value of 2
 - C. Based on the port of 50
 - D. Based on the port of 51
7. Which of the following choices provides the best protection against potentially malicious FTP commands?
 - A. Defense diversity
 - B. Packet-filtering firewall
 - C. Stateful inspection firewall
 - D. Application firewall
8. How can you provide defense diversity with a DMZ?
 - A. Use a single firewall.
 - B. Use two firewalls from the same vendor.
 - C. Use two firewalls from different vendors.
 - D. Ensure that only trusted partners are allowed access.
9. It's common to enable or install a firewall on a server to protect the server. What type of firewall is this?
 - A. Network-based
 - B. Hardware-based
 - C. Packet-filtering
 - D. Host-based
10. Of the following choices, what is NOT used for VPNs?
 - A. L2TP
 - B. PPTP
 - C. SSLTP
 - D. TLS

11. What port does PPTP typically use?
 - A. 143
 - B. 443
 - C. 1701
 - D. 1723
12. What port does a TLS VPN typically use?
 - A. 80
 - B. 88
 - C. 143
 - D. 443
13. How would users typically access a TLS VPN?
 - A. With a web browser
 - B. With a dedicated application
 - C. With broadband access but never DSL access
 - D. With an IMAP application
14. Of the following choices, what indicates the primary improvement that MS-CHAPv2 included over previous protocols?
 - A. Support for biometrics
 - B. Use of certificates
 - C. Mutual authentication
 - D. Use of a nonce
15. Which of the following identifies the correct representation of RADIUS?
 - A. Remote Access Dial-in User System.
 - B. Remote Authentication Dial-in User Service
 - C. Roaming Access Dial-in User Service
 - D. Remote Authentication Dialing User System
16. What port does TACACS+ typically use?
 - A. 25
 - B. 49
 - C. 53
 - D. 443

17. What can be used to examine the health of a client prior to allowing network access and restricting access of unhealthy clients to a quarantined network?
 - A. RADIUS
 - B. TACACS+
 - C. NAC
 - D. SRTP
18. Which of the following represents the greatest risk to virtual systems?
 - A. Confidentiality
 - B. VM escape
 - C. Increased costs for power and cooling
 - D. Loss of control of data in the cloud
19. Which of the following is an example of SaaS?
 - A. Access to an operating system over the Internet
 - B. Access to a server over the Internet
 - C. Web-based e-mail
 - D. VM escape
20. An organization is sharing resources with another organization using cloud-based computing. Which of the following cloud operation models does this describe?
 - A. Community
 - B. Hybrid
 - C. Private
 - D. Public

Answers

1. **B.** A table in a switch maps the media access control (MAC) address to the physical port. Routers map network IP addresses to the physical port with the corresponding gateway IP addresses. The Address Resolution Protocol (ARP) resolves IP addresses to MAC addresses, but this data is not held in a switch. Physical ports and logical ports are not mapped together.
2. **A.** A virtual local area network (VLAN) segments traffic on a network using a switch. Extensible Authentication Protocol (EAP) is used for authentication, not segmentation. Secure Sockets Layer (SSL) and Transport Layer Security (TLS) are transport encryption protocols and do not segment traffic.

3. **D.** The Secure Real-time Transport Protocol (SRTP) provides confidentiality, authentication, and replay protection for Voice over IP (VoIP) transmissions. Remote Authentication Dial-in User Service (RADIUS) and Terminal Access Controller Access Control System+ (TACACS+) are used to provide authentication, authorization, and accounting (AAA) for remote access. The public switched telephone network (PSTN) is one of the methods used for Internet access.
4. **B.** The primary way to control call forwarding is to restrict numbers that can be used for call forwarding. Protecting the administrator password and changing it often protects the overall system, but doesn't directly address call forwarding. Restricting long distance calling is also important, but it doesn't address call forwarding. Although physical security of the phone system is valuable, it won't control call forwarding.
5. **A.** A proxy server can cache web pages that are retrieved from the Internet. It can also block users from accessing restricted websites by filtering the web page requests. A proxy server does not provide authentication directly, although some proxy servers can be tied into an authentication system. A proxy server does not normally perform firewall functions.
6. **A.** Packet-filtering firewalls can filter traffic based on IP addresses, ports, and protocol IDs, and a protocol ID of 1 identifies Internet Control Message Protocol (ICMP) traffic. Internet Group Message Protocol (IGMP) uses a protocol ID of 2. Internet Protocol security (IPsec) is not identified by ports. IPsec Encapsulating Security Protocol (ESP) has a protocol ID of 50, and IPsec Authentication Header (AH) has a protocol ID of 51.
7. **D.** An application firewall (also called an application proxy or an application gateway firewall) can inspect commands used by individual protocols such as File Transfer Protocol (FTP) and block potentially malicious commands. Defense diversity refers to using firewalls from two different vendors in a demilitarized zone (DMZ). A packet-filtering firewall can only inspect individual packets for IP addresses, ports, and protocol IDs. A stateful inspection firewall can track the activity within TCP and UDP sessions, but can't interpret commands.
8. **C.** You can provide defense diversity with a DMZ by using two firewalls from different vendors. If a vulnerability appears in one, it's unlikely that a vulnerability will exist in the second firewall at the same time (unless the second is from the same vendor). A single firewall doesn't provide any diversity. An extranet (not a DMZ) would allow access only to trusted partners.
9. **D.** A host-based firewall is installed or enabled on individual hosts, such as desktop computers or servers, and provides protection for the host. Network-based firewalls protect the network rather than individual systems. Packet filtering identifies the method used by the firewall, and both network-based and host-based firewalls can filter packets.

10. **C.** There is no such thing as SSLTP in the context of virtual private networks (VPNs). The other choices (L2TP, PPTP, and TLS) are used for VPNs.
11. **D.** The Point-to-Point Tunneling Protocol (PPTP) uses TCP port 1723. Internet Message Access Protocol version 4 (IMAP4) uses TCP port 143, Transport Layer Security (TLS) and Secure Sockets Layer (SSL) use TCP port 443, and Layer 2 Tunneling Protocol (L2TP) uses UDP port 1701.
12. **D.** A Transport Layer Security (TLS) virtual private network (VPN) typically uses TCP port 443, the same port as HyperText Transfer Protocol Secure (HTTPS). HyperText Transfer Protocol (HTTP) uses TCP port 80. Kerberos uses TCP port 88. Internet Message Access Protocol version 4 (IMAP4) uses port 143.
13. **A.** Users typically access a Transport Layer Security (TLS) virtual private network (VPN) using a web browser instead of a dedicated application. A TLS VPN is not dependent on a specific type of Internet connection (such as broadband or DSL). Internet Message Access Protocol (IMAP) is used with e-mail, not VPNs.
14. **C.** MS-CHAPv2 uses mutual authentication, where the client authenticates to the server and the server authenticates to the client. MS-CHAPv2 does not directly support biometrics. MS-CHAPv2 can be used with the Extensible Authentication Protocol (EAP) to support certificates, but it cannot do so on its own. CHAP uses a nonce (a number used once), so this isn't an improvement.
15. **B.** RADIUS is an acronym for Remote Authentication Dial-in User Service. The other choices are not valid.
16. **B.** Terminal Access Controller Access Control System+ uses TCP port 49. Simple Mail Transport Protocol (SMTP) uses TCP port 25, Domain Name System (DNS) uses TCP port 53 and UDP port 53, and HyperText Transfer Protocol Secure (HTTPS) uses TCP port 443.
17. **C.** A network access control (NAC) system can check a system's health based on a predefined health policy and restrict the access of unhealthy clients to a quarantined network. Remote Authentication Dial-in User Service (RADIUS) and Terminal Access Controller Access Control System+ (TACACS+) are used to provide authentication, authorization, and accounting (AAA) for remote access. Secure Real-time Transport Protocol (SRTP) provides confidentiality, authentication, and replay protection for Voice over IP (VoIP) transmissions.
18. **B.** VM escape is a known attack against virtual systems. If the attack is successful, an attacker can access the host system and all virtual systems within the host. Loss of confidentiality (not confidentiality) is a risk that can be reduced with encryption. Virtualization reduces costs for power and cooling. Loss of control of data stored in the cloud is a risk associated with cloud computing, but organizations can use virtual systems internally to keep control of their data.

- 19. C.** Web-based e-mail is an example of Software-as-a-Service (SaaS). SaaS, also known as on-demand software, provides users with access to software or applications over the Internet. Platform-as-a-Service (PaaS) is a cloud computing service where users have access to a platform with an operating system. Infrastructure-as-a-Service (IaaS) provides users with access to hardware such as servers or network devices. VM escape is an attack on virtual systems.
- 20. A.** A community cloud is a private cloud that is shared by two or more organizations. A hybrid cloud is a combination of any two or more clouds. A private cloud is only available to users within an organization. Public cloud-based services are provided by third-party vendors and are available to anyone.

Attacks

In this chapter, you will learn about

- Differences between hackers and crackers
- Differences between white hats, black hats, and gray hats
- Advanced persistent threats
- Basic countermeasures used to thwart attacks
- Common denial of service (DoS) and distributed DoS (DDoS) attacks
- Botnets and zombies
- Software security countermeasures
- Phishing, spear phishing, and whaling
- Social engineering tactics and countermeasures
- The importance of training as a countermeasure

Comparing Attackers

Attackers break into computers and networks with malicious intent and often for personal gain. Sometimes, they try to take down a system or put it out of service. Other times, they try to get into the system to gain information. For example, an attacker may break into a system to steal credit card data from a business's customers. The attacker can then use the credit cards fraudulently. Alternatively, they may attempt to bring a system down as an act of revenge against an organization.

Make no mistake, though. These attackers are criminals. Just as someone who breaks into your home is a criminal, attackers who break into networks and individual systems are also criminals. They are breaking laws. Many attack individuals and networks for monetary gain without any apparent care to how the loss affects their victims.

In general, criminals must have motive, means, and opportunity to complete a crime. The motive for an attacker is often money, but can also be revenge, espionage, or other reasons. The means is often with some type of malicious code, and the opportunity is available for anyone with Internet access. Indeed, for a greedy criminal willing to risk jail time, there are often enough programs available on the Internet that they don't even have to write the malicious code, but only need to locate and download it. The following sections identify the different types of attackers and some names used to identify them.

Hackers and Crackers

Two terms you'll come across when studying IT security are hackers and crackers. The following two bullets provide traditional definitions (even though many people do not use these traditional definitions):

- **Hacker** Someone that is proficient with computers (often with programming skills) and uses these skills to gain and share knowledge with others. By this definition, a hacker does not break into systems with malicious intent or for personal gain.
- **Cracker** Someone that is proficient with computers and uses these skills to attack systems. A cracker *does* attack systems with malicious intent or for personal gain.

The English language has a long history of changing the use of words and their meanings, often with some people kicking and screaming along the way. For example, you probably don't use *thee* and *thou* in your daily speech. However, there was a time when everyone used these words regularly.

Similarly, the word *hacker* is evolving. Currently, the media widely uses the term *hacker* to identify anyone launching attacks on computers or networks. Even though some IT people are clinging to the traditional definition, it has lost the traditional meaning outside of these circles.



NOTE Based on the traditional definitions, hackers are good, and crackers are evil. However, the media commonly uses the term *hacker* to describe anyone who attacks with malicious intent and rarely uses the word *cracker* in this context. In this book, I'm using the term *attacker* to refer to anyone who attacks a network.

White Hats, Black Hats, and Grey Hats

Using the traditional definitions, hackers are good and crackers are evil. However, since the media began confusing these terms, the concept of white hats, black hats, and gray hats has appeared. These designations are reminiscent of the older cowboy movies where the good guys wore white hats and the bad guys wore black hats.

White hats fall into the traditional definition of hackers and are sometimes called ethical hackers. Organizations often hire them as security consultants to perform vulnerability assessments or provide other security services.

In contrast, black hats fall into the traditional definition of a cracker. Black hats are criminals who attack systems with malicious intent or for personal gain.

Of course, good and evil don't always have definitive lines with good on one side and evil on the other. There are gray areas. Similarly, gray hats are individuals who have exceptional computer and networking skills, but they don't use them for personal gain or with malicious intentions. However, their activities may cross ethical boundaries.

As a comparison between the three, imagine that someone discovers a weakness or vulnerability in a system. A white hat would inform the owners of the system but not take any further action. A black hat would try to exploit the vulnerability without telling the owner. A gray hat may first try to let the owners know of the vulnerability, but then decide to publish the vulnerability if the gray hat doesn't think the owner is acting responsibly. By publishing the vulnerability, the gray hat makes the vulnerability known to the black hats, who then use the information to exploit it.

Advanced Persistent Threats

An advanced persistent threat (APT) refers to a highly sophisticated group of people who target their attacks. They have both the capability and the intent to carry out successful attacks and are typically sponsored by a government. Once they identify a target or goal, they persistently attack until they succeed. Many countries reportedly have teams of IT personnel regularly probing the IT systems of organizations and other government networks.

As an example, in 2013 security firm Mandiant, now a division of FireEye, released a report titled *APT1, Exposing One of China's Cyber Espionage Units* (available for download at <http://intelreport.mandiant.com/>). The report documents the activities of a Chinese group that Mandiant named "APT1." Mandiant is very clear in its conclusions and the report includes this statement: "The details we have analyzed during hundreds of investigations convince us that the groups conducting these activities are based primarily in China and that the Chinese Government is aware of them." The report documents the following activities by APT1:

- Compromised at least 141 companies in 20 different industries in many different countries, including the United States, Canada, and the United Kingdom
- Maintained access in compromised systems for an average of 365 days and as long as 58 months
- Stole a large quantity of intellectual property, including 6.5 terabytes of data from a single organization
- Developed over 40 different families of malicious software (malware)

Mandiant reported that the most commonly observed method APT1 used to gain access to a network was a spear phishing attack (described later in this chapter). This type of attack establishes a foothold by infecting a single system, and then the attackers quickly infect other systems in the network. For example, they often install backdoors on internal systems, providing the attackers with remote access to the systems.

An important point to remember here is that APT1 is only one APT operating out of China, and there may be more. Similarly, other countries very likely have their own APTs. APTs don't attempt to get all the information immediately. Instead, they are patient and willing to get a little bit of information at a time. However, they are also persistent. When one attack fails, they don't give up. Instead, they look for other ways to attack and keep pursuing the target until they achieve some success.

After many successful attacks against computers and networks in the United States, the U.S. government created the United States Cyber Command (USCYBERCOM), which became operational in 2010. Its mission statement is as follows:

USCYBERCOM plans, coordinates, integrates, synchronizes, and conducts activities to: direct the operations and defense of specified Department of Defense information networks and; prepare to, and when directed, conduct full-spectrum military cyberspace operations in order to enable actions in all domains, ensure US/Allied freedom of action in cyberspace and deny the same to our adversaries.

Many other countries have created similar organizations to combat the threats of cyberwarfare.

Insider Attacks

An insider is someone who works within an organization, such as an employee, consultant, or vendor, who has authorized access to the organization's assets. Insiders have the means and opportunity to cause a significant amount of damage and can cause the loss of confidentiality, integrity, and availability of an organization's resources.

Insiders may steal data such as trade secrets, customer information, or other classified information from the organization. They can then sell this information for personal gain.

Some insiders are out for revenge. For example, Rajendrasinh Makwana was terminated from a contractor job at Fannie Mae in 2008. However, he retained access to the systems for at least a short period afterward. He later embedded a time bomb script in the Fannie Mae system. It was set to run on January 31, 2009, and would have deleted passwords on about 4,000 servers, erased all the data and backup data for the servers, and then powered them down. He also programmed the script to disable the ability to turn the servers back on remotely. Had the script run, the result would have been catastrophic. Luckily, another engineer discovered the damaging script before it ran. Interestingly, Makwana's employer terminated him for allegedly making unauthorized changes to other systems. However, the employer didn't revoke his access right away, which gave him enough time to install the malicious script. A jury convicted Makwana and he was ultimately sentenced to over three years in prison.

Although insider attacks have been the primary threat to computer systems in the past, this trend has reversed. Today, the primary threats come from attackers outside the organization, and these attacks are typically more costly than insider attacks. Based on several different reports, attacks from insiders have been a significantly smaller threat than attacks from outsiders. As an example, the annual Verizon *Data Breach Investigations Report* (years 2008 through 2015) has consistently shown that external sources are responsible for the majority of data breaches. The 2015 report indicated insider issues were only about 10 percent of the total number of attacks.

Even though insider attacks continue to decline, it's worth stressing that an organization must protect assets against attacks from both outsiders and insiders. Effective access controls (including strong authentication and authorization mechanisms) combined with strong auditing techniques help to prevent successful attacks from insiders.

Some organizations are using sophisticated monitoring systems to detect and prevent insider attacks before they can cause damage. For example, in 2009 the Transportation Security Administration's (TSA) Colorado Springs Operations Center internal surveillance system identified the actions of Douglas Duchak as he was planting malware in a critical system about a week after he was terminated with a two-week notice. He was sentenced to a two-year prison sentence.

Script Kiddies

A script kiddie is a relatively inexperienced attacker who can run a script or application to launch an attack. The script kiddie rarely has the knowledge to program the script or application and often doesn't understand the details of what the attack is doing. You can think of a script kiddie as a bored teenager who downloaded some hacking tools and used them to launch attacks.

Insiders vs. APTs

Years ago, many security professionals considered insider attacks to be the primary threat to computer systems. However, this has changed as advanced persistent threats (APTs) and criminal groups have become more active.

As an example, an organized group calling itself GOP (short for Guardians of Peace) attacked Sony Pictures Entertainment in late 2014. The group stole over 25GB of data. It then began leaking information about more than 4,000 past and present employees. This resulted in losses to Sony as high as \$300 million. This was in addition to the \$170 million in losses after the 2011 Sony PlayStation attack. The U.S. government ultimately imposed sanctions on North Korea after the Sony attack, blaming the North Korean government for the attack. This indicates an APT operating out of North Korea launched the 2014 attacks on Sony.

As another example, an organized group of criminals operating out of Russia stole more than \$100 million from people and organizations. These criminals primarily used two attacks: GameOver Zeus and CryptoLocker. GameOver Zeus infects a system and steals banking credentials, which the criminals use to empty bank accounts and move funds to offshore accounts. CryptoLocker encrypts the data on a system and demands ransom money to decrypt the data.

Both GameOver Zeus and CryptoLocker include additional malicious actions. However, the important point here is that an organized group of criminals controls the attacks, and these criminals are much more of a threat than insiders are.

This isn't to say that an organization shouldn't protect itself against malicious insiders. However, many common security practices (such as implementing the principle of least privilege and separation of duties policies) limit the potential losses against malicious insiders.

It's worth stressing that although script kiddies exist, many attackers today are highly proficient at programming techniques. They have detailed knowledge of different attack methods and are constantly modifying existing attacks to prevent detection.

Phreaks

A phreak (also known as a phone phreak or phreaker) is someone who illegally breaks into a phone system. The phreaker's goal may be to access the system to make long distance calls or to tap the phone lines. Phone phreaking is the practice of hacking into a phone system.

Accidental Threats

Although the infamous disgruntled worker can certainly cause a significant amount of damage intentionally, employees can also cause damage accidentally. Some examples include the following:



NOTE Accidental threats aren't attacks. Employees who accidentally cause damage aren't doing so with malicious intent or for personal gain. However, some incidents caused by accidental threats can cause significant damage, so it's worthwhile knowing about them.

- **Responding to phishing attempts** Phishing is explained in more depth later in this chapter, but in short, it's an e-mail that encourages the user to respond with personal information or to click a link. The user can give up valuable company information or inadvertently install malicious software on their system. This provides a method for an outside attacker to attack the organization.
- **Forwarding malicious software (malware) or bringing malware from home** An uneducated user can forward viruses or other forms of malware via e-mail or bring it on a USB drive from a home computer. The malware could cause damage to internal systems or use other methods to gather information for attacks against the organization.
- **Unauthorized data access** If adequate access controls aren't in place, a user may be able to access data. The user may then accidentally modify or delete it, or even disclose the contents of the data to someone else without realizing its true value. Something as simple as an unauthorized employee gaining access to personal records of other employees can cause a great deal of havoc to an organization.
- **Losing hardware** Many organizations issue hardware such as laptops or USB drives to users. When users don't recognize the value of the hardware and the data it holds, they may not provide adequate protection, resulting in its loss. As an example, many users have left laptops in conference rooms during a lunch break, only to discover them missing after lunch.

Exploring Attack Types and Countermeasures

There are numerous attack types and countless variations. As attackers have launched new types of attacks, IT security professionals have identified new methods of protection. When one attack method no longer works, attackers develop others. It's a constant battle between the attackers and security professionals. As long as attackers are able to achieve their objectives through attacks, expect them to continue.

This section outlines many of the common types of attacks and some countermeasures. This is by no means a complete list, but it does provide an overview.

Basic Countermeasures

Countermeasures are security controls that reduce risks. They typically reduce vulnerabilities or reduce the impact of a threat. Chapter 9 covers countermeasures and controls in more depth. The following sections introduce some basic countermeasures that help prevent many attacks.

Patching Systems

Many attacks are possible because of vulnerabilities discovered in an operating system. However, once vendors learn about these vulnerabilities, they develop, test, and release patches and updates to plug the security holes. As long as administrators keep systems up to date with current patches and updates, the systems are no longer vulnerable to that specific attack. However, unpatched systems remain vulnerable to the attacks.

Hardening Systems

Administrators harden systems by increasing the security provided with the default configuration. Hardening includes removing or disabling unneeded protocols and services, ensuring the system is deployed with all current patches and updates, and enabling software firewalls.

Increasing User Awareness

User awareness programs help users understand risks, recognize risky behaviors, and understand their responsibilities related to security. A goal is to reaffirm safe computing practices and encourage users to change their behavior if they are engaging in unsafe computing practices.

Spoofing

Spoofing refers to impersonating or masquerading as someone or something else, and many attacks use spoofing techniques. Chapter 3 presented the concept of spoofing in the context of IP addresses and media access control (MAC) addresses. Attackers can change the source IP address or source MAC address so that traffic appears to be coming from a different computer. Similarly, attackers can change the From field in an e-mail to make it appear as though the e-mail is coming from someone else.

DoS

A denial of service (DoS) attack attempts to prevent a system from answering legitimate requests from users. The attack often attempts to overwhelm the system by consuming system resources. While the attacked system responds to the attack, it doesn't have enough resources to respond to regular requests. The goal is to deprive the organization or its customers from the use of a service.



EXAM TIP A denial of service (DoS) attack attempts to disrupt a service. It can do so by taking a system down or by consuming system resources with meaningless requests that overload its ability to respond to legitimate requests. The effect is that valid authorized users are unable to access a system.

As an example, a common DoS attack is a SYN flood attack that initiates the TCP handshake process but doesn't complete it. Chapter 3 describes the three-way handshake used by TCP. Recall that a normal three-way handshake includes the following three packets:

1. A TCP client sends a packet with the SYN flag set to a system such as a server.
2. The server responds with a packet with the SYN and ACK flags set.
3. The client then completes the handshake process by sending back a packet with the ACK flag set.

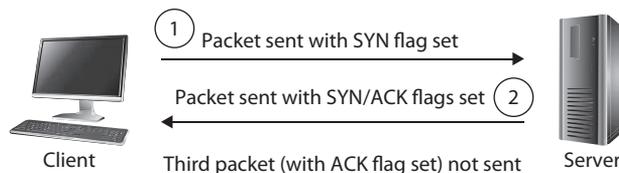
In a SYN flood attack, the attacker withholds the third packet, leaving the half-open session on the server. For example, consider Figure 5-1, which shows the process of a SYN flood attack. The first two packets are sent just like a regular three-way handshake, but the third packet is not sent. Instead, the attacker then sends out another SYN packet.



EXAM TIP In a SYN flood attack (also called a TCP SYN, TCP flood, and TCP half-open attack), the attacker floods a system with SYN packets but withholds the third packet in the TCP handshake process.

This is similar to a practical joker extending his hand to shake yours, you extending your hand in response, and the joker pulling away his hand at the last second. He sticks his hand out again, you extend your hand again, and he pulls his hand away again. You probably won't fall for this repeatedly. However, TCP will keep responding to SYN packets with SYN/ACK packets and continue waiting for the third packet.

Figure 5-1
SYN flood attack





TIP Most IDSs can detect this type of attack. If the attack is coming from a single IP address, an IDS can change the access control list (ACL) on the firewall or a boundary router to block all traffic from the offending IP address.

If this happens once or twice, it's no big deal. However, an attacker can initiate hundreds of half-open sessions in a short period. If the attack isn't detected, the server can keep each of these sessions open as it waits for hundreds of ACK packets. Each one of these sessions consumes resources. The attack can slow down the server, and in extreme cases, it can actually crash the server.

The following are some other well-known DoS attacks:

- **Ping of death** An Internet Control Message Protocol (ICMP) ping packet is normally 32 bytes of data, although different operating systems may vary the size. A ping of death attack changes the size of the ping packet to more than 64KB. These large ping packets have crashed some systems, although most systems recognize the attack today and can avoid it.
- **LAND** A local area network denial (LAND) attack tricks a system into sending out packets to itself in an endless loop. In other words, the source and destination addresses are set to the attacked system. A LAND attack can shut down a system.



TIP Although many of these attacks will fail on protected systems, they are valuable to understand. Many current attacks are simply variations on attacks that have been successful in the past.

DDoS

A distributed denial of service (DDoS) attack is similar to a DoS attack except that it comes from multiple attackers at the same time. For example, an attacker may launch a TCP flood attack against a single system from hundreds or even thousands of systems. A key to the success of DDoS attacks is for the attacker to take control of multiple systems and then use these systems to launch the attacks. Attackers often use botnets to launch DDoS attacks.



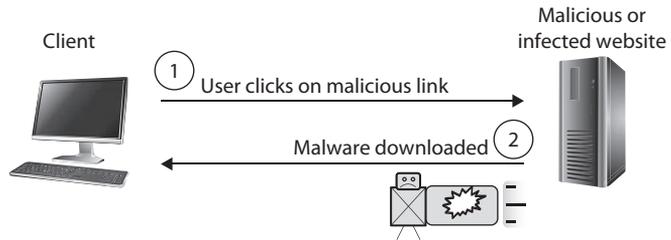
EXAM TIP The first step in a DDoS attack is often to compromise as many systems as possible. The attacker then directs these systems to launch the DDoS attack.

Botnets and Zombies

A botnet is a group of computers (often called *zombies*) controlled by an attacker. The term *botnet* is a combination of *robot* and *network*. The attacker (sometimes called a bot herder) manages a command and control center and directs the computers in the botnet to do the bidding of the attacker.

Computer users are usually unaware that their computer has become a zombie. Instead, malware infects computers without the user's knowledge. This allows attackers to control the computer from remote locations.

Figure 5-2
Single system
infected in drive-
by download



One way that a system is infected and joined to a botnet is through a drive-by download, as shown in Figure 5-2. When a user visits an infected or malicious website, the website downloads the malware and attempts to install it on the user's system. Users sometimes unknowingly visit these sites during Internet surfing sessions. Other times, users click links included within malicious e-mails. As an example, the GameOver Zeus malware used drive-by downloads to join computers to botnets. Additionally, the attackers used malicious e-mails to encourage users to visit the malicious websites.



TIP Drive-by downloads are any downloads that occur without the user's knowledge after the user visits a website. They can include many different types of malware, including software that joins the computer to a botnet.

While a drive-by download is one way that computer systems are joined to botnets, it isn't the only way. Multiple types of malware can include code to join computers to botnets. Chapter 6 covers malware in more depth.

Figure 5-3 shows a botnet with multiple zombies in it. These zombies regularly check in with the attacker's command and control center, which is hosted on one or more

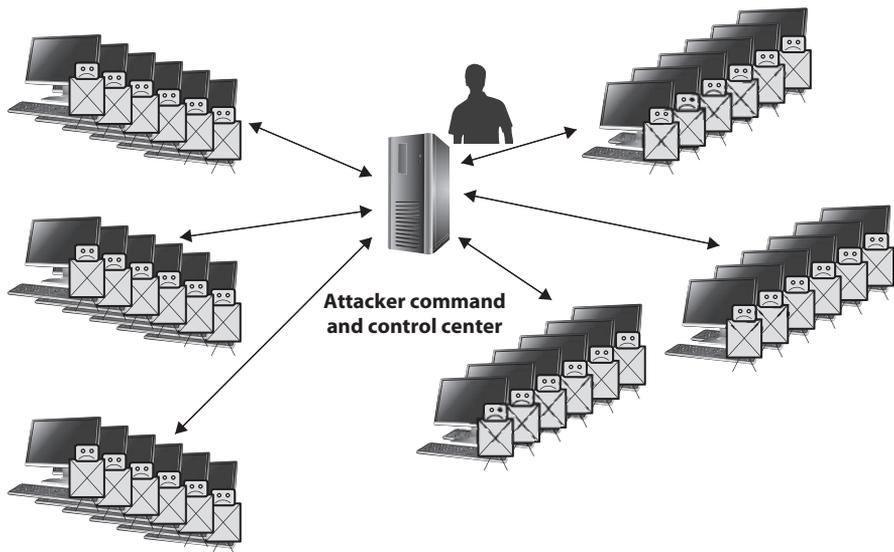


Figure 5-3 Zombies regularly check in with the command and control center.

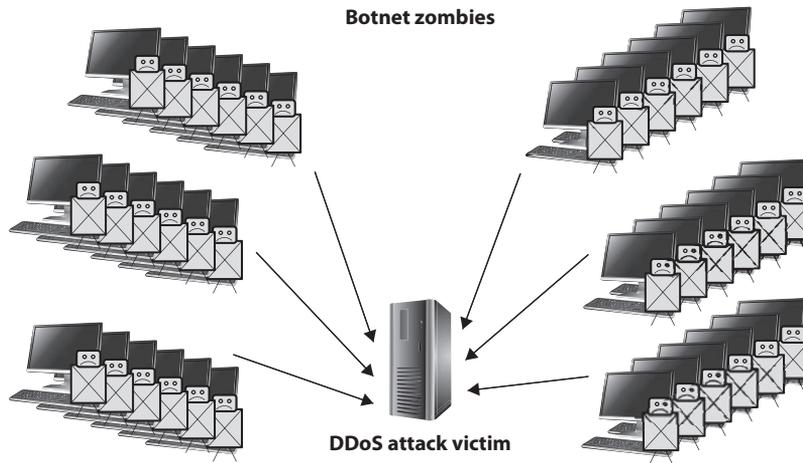


Figure 5-4 Zombies can launch DDoS attacks.

servers controlled by the bot herder. Smaller botnets may have a single server hosting the command and control software. Larger botnets may have dozens of servers located around the world.

The commands issued to the zombies in the botnet vary. Bot herders can direct zombies to bombard a server with a DDoS attack, as shown in Figure 5-4. Attackers can also direct zombies to take other malicious actions, such as send spam and phishing e-mails, spread malware, or launch other types of attacks.

Criminals often rent out their botnets to other criminals, and if the botnet is large, the owners can generate a significant amount of money from these rentals. Many botnets control tens of thousands of computers, and some control more than a million. At its height, the Rustok botnet had more than a million zombies and sent out billions (yes, billions with a *b*) of spam e-mails daily. The Rustok botnet was taken down in March 2011, and IT watchers reported that spam levels immediately plummeted.



NOTE Microsoft launched the Microsoft Digital Crimes Unit (DCU) in 2010 and the Microsoft Cybercrime Center in 2013. They both work with other organizations to fight cybercrime, and have helped take down multiple botnets around the world.

Sniffing Attack

A protocol analyzer (also called a sniffer or packet sniffer) captures packets transmitted over a network. After capturing packets, you can save them in a file (commonly called a capture) and browse them at your leisure. IT professionals use sniffers to troubleshoot and learn details about a network. Attackers use sniffers to capture data and exploit it for personal gain.

A popular sniffer is Wireshark. It's available as a free download and runs on most platforms, such as Windows, UNIX, and Linux.



NOTE You can read more about Wireshark and download a copy from here: www.wireshark.org. Wireshark was previously known as Ethereal.

If data is sent over the network in cleartext, an attacker can easily read it in a sniffer. However, if the data is encrypted, it can't be read. For example, Chapter 3 mentioned the risks of unencrypted protocols such as File Transfer Protocol (FTP) and Trivial File Transfer Protocol (TFTP). Both FTP and TFTP send data over the network in cleartext by default. An attacker can capture these transmissions and read credential information, such as usernames and passwords, and the actual data. In contrast, FTP Secure (FTPS) uses Secure Shell (SSH) to encrypt the transmissions, preventing attackers from reading data.

Figure 5-5 shows a screenshot of Wireshark with a capture. During this capture, I connected to a network drive and opened a file named passwords.txt. After stopping the capture, I browsed through it to locate packets that held data. In the figure, you can see arrow 1 pointing to packet 1470, which is an SMB Read AndX Response command. The

The screenshot shows the Wireshark interface with the following details:

- Filter:** Expression... Clear Apply
- Packet List:**

No.	Time	Source	Destination	Protocol	Info
1462	07.111991	192.168.1.134	192.168.1.141	SMB	Trans2 Response, QUERY_FILE_INFO
1463	87.112107	192.168.1.141	192.168.1.134	SMB	Trans2 Request, QUERY_FILE_INFO, FID: 0x16f5, Query File Basic Info, Path: \\mybookworld\public\tmp\passwords.txt
1464	87.113164	192.168.1.134	192.168.1.141	SMB	Trans2 Response, FID: 0x16f5, QUERY_FILE_INFO
1465	87.115689	192.168.1.141	192.168.1.134	SMB	Trans2 Request, QUERY_PATH_INFO, Query File Basic Info, Path: \\mybookworld\public\tmp\passwords.txt
1466	87.117888	192.168.1.134	192.168.1.141	SMB	Trans2 Response, QUERY_PATH_INFO
1467	87.118132	192.168.1.141	192.168.1.134	SMB	Trans2 Request, QUERY_PATH_INFO, Query File Standard Info, Path: \\mybookworld\public\tmp\passwords.txt
1468	87.120342	192.168.1.134	192.168.1.141	SMB	Trans2 Response, QUERY_PATH_INFO
1469	87.122470	192.168.1.141	192.168.1.134	SMB	Read AndX Request, FID: 0x16f5, 26 bytes at offset 0
1470	87.123714	192.168.1.134	192.168.1.141	SMB	Read AndX Response, FID: 0x16f5, 26 bytes
1471	87.124647	192.168.1.141	192.168.1.134	SMB	Trans2 Request, QUERY_PATH_INFO, query File Basic Info, Path: \\mybookworld\public\tmp\passwords.txt
1472	87.126384	192.168.1.134	192.168.1.141	SMB	Trans2 Response, QUERY_PATH_INFO
1473	87.126895	192.168.1.141	192.168.1.134	SMB	Trans2 Request, QUERY_PATH_INFO, Query File Basic Info, Path: \\mybookworld\public\tmp\passwords.txt
- Packet Details:**
 - Frame 1470: 143 bytes on wire (1144 bits), 143 bytes captured (1144 bits)
 - Ethernet II, Src: WesternD_82:4b:5f (00:90:a9:82:4b:5f), Dst: Microsoft_a6:73:6c (00:03:ff:a6:73:6c)
 - Internet Protocol, Src: 192.168.1.134 (192.168.1.134), Dst: 192.168.1.141 (192.168.1.141)
 - Transmission Control Protocol, Src Port: netbios-ssn (139), Dst Port: 52466 (52466), Seq: 623497, Ack: 42178, Len: 89
 - NetBIOS Session Service
 - SMB (Server Message Block Protocol)
 - SMB Header
 - Read AndX Response (0x2e)
 - [FID: 0x16f5 (\mybookworld\public\tmp\passwords.txt)] ← 2
 - word count (wct): 12
 - AndXCommand: No further commands (0xff)
 - Reserved: 00
 - AndXOffset: 0
 - [File offset: 0]
- Packet Bytes:**

```

0000  00 03 ff a6 73 6c 00 90 a9 82 4b 5f 08 00 45 00  ...s!.. ..K...E.
0010  00 81 a7 03 40 00 40 06 0f 10 c0 a8 01 86 c0 a8  ...@.@.....
0020  01 8d 00 8b cc f2 74 15 9c 14 4e 2a e5 89 50 18  ...t... ..N*..P.
0030  18 70 8e 3a 00 00 00 00 00 55 ff 53 4d 42 2e 00  ...p.....U.SMB..
0040  00 00 00 88 01 c8 00 00 00 00 00 00 00 00 00 00  ...e...DZ....
0050  00 00 02 00 ff fe 65 00 44 5a 0c ff 00 00 00 ff  ...Dar rll G!Bs
0060  ff 00 00 00 00 1a 00 3b 00 00 00 00 00 00 00 00  ...on..!Pa3 tedScR
0070  00 00 00 1a 00 44 61 72 72 69 6c 20 47 69 62 73
0080  6f 6e 0d 0a 49 30 61 24 24 65 64 53 53 43 50

```

Figure 5-5 Viewing text data in Wireshark

center pane holds the details on this packet. Arrow 2 shows where the packet includes the name of the file (passwords.txt). The bottom pane shows the data in both a hexadecimal and text format, and arrow 3 shows the username and password in cleartext (Darril Gibson, IPa\$\$edSSCP).

Sniffers can work in either promiscuous mode or nonpromiscuous mode. In nonpromiscuous mode, the sniffer will capture only data sent directly to or from the computer's IP address. It will not capture any packets that do not have the computer's IP address as the source or destination address. In contrast, when the sniffer works in promiscuous mode, it will capture all data that reaches the computer regardless of the IP address. Technically, the sniffer configures the network interface card (NIC) to work in promiscuous mode for the application.



EXAM TIP Promiscuous mode allows a sniffer to capture all data that reaches the sniffer. When the sniffer runs in nonpromiscuous mode, it will only capture data sent directly to or from the system running the sniffer.

One of the ways an organization reduces vulnerabilities from sniffer attacks is to use switches instead of hubs. For example, consider Figure 5-6, which shows both a hub and a switch in action. In the figure, PC1 and PC3 are sending data to and receiving data from each other. The hub sends the data to all ports, so an attacker with a sniffer running on PC2 is able to capture the transmission with a sniffer (as long as the sniffer is running in promiscuous mode).

However, when using a switch, the same data transmissions reach only PC1 and PC3. The attacker running a sniffer from PC2 is not able to capture the data because the traffic doesn't ever reach the switch port for PC2.

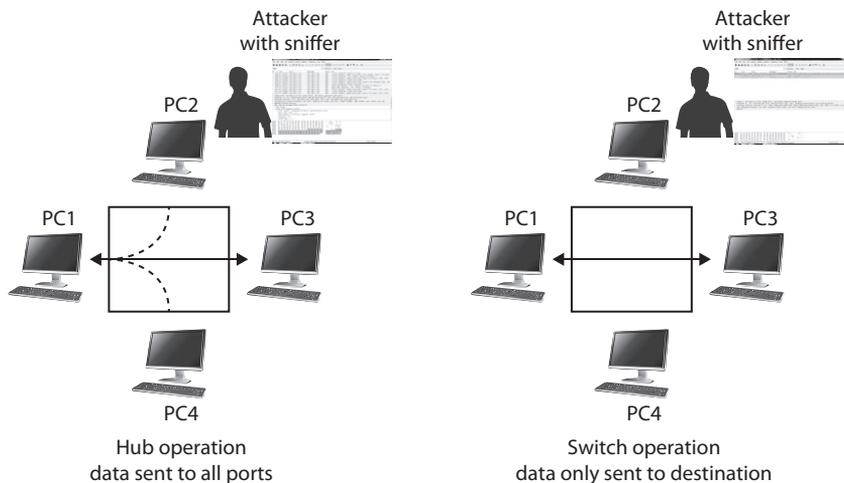


Figure 5-6 Comparing data sent via a hub and via a switch

It's worthwhile noting that administrators can configure a mirrored port on most routers and switches. The device sends a copy of all data sent through the device to the mirrored port. If attackers can connect directly to the mirrored port, they will be able to capture all the traffic going through the device. This underscores the importance of physical security. Locating network devices in secure server rooms or locked wiring closets provides a necessary layer of security.

Ping Sweep

Attackers often start an attack with a ping sweep. A ping sweep uses ICMP to identify what systems are operational within a range of IP addresses. This is similar to sending a ping to a system and waiting for a response. However, instead of sending a single ping to a single system, the ping sweep sends pings to all the systems within a range of IP addresses.

A ping sweep can be used as part of an overall attack strategy. The ping sweep performs reconnaissance to identify the system. Other tools (such as port scanners or vulnerability scanners) identify details on individual systems in a fingerprinting attack.

Port Scan

The port scan attempts to detect what ports are open on a system. Attackers often use it as part of an overall fingerprinting attack to gather information on a system, such as what ports are open. Chapter 3 covered well-known ports. As a reminder, well-known ports are in the range of 0 to 1023 and are commonly used by specific protocols. For example, SSH uses port 22, Simple Mail Transfer Protocol (SMTP) uses port 25, and HyperText Transfer Protocol (HTTP) uses port 80.



EXAM TIP A fingerprinting attack attempts to gather information on specific computers, such as what protocols are running on the system or what specific functions the server performs. In contrast, a reconnaissance attack is a broad-based attack that attempts to identify systems on a network. For example, a ping sweep can identify computers based on their response to ping requests. A port scan on these computers provides more information on them.

If a port scan determines that port 80 is open, it indicates that HTTP is enabled on a system. Web servers use HTTP, so this also indicates the system is probably a web server. It could be an Apache web server running on a Linux system, an Internet Information Services (IIS) web server running on a Microsoft system, or something else. Attackers often use additional steps to identify the server's operating system and the web server software.

Port scanning tools can run sequentially through ports. You can detect them by checking firewall logs or by using automated IDSs. For example, if the same source IP address checks for responses on ports 20, then 21, then 22, and so on, it's a clear indication of a port scan. More sophisticated port scanners can randomize the order of the port numbers that they check, such as 22, then 80, then 25, and so on. However, most IDSs are aware of both tactics and can detect both sequential and randomized port scans.

Salami Attack

In a salami attack, an attacker performs a large number of minor actions that likely won't be noticed or reported but collectively can add up to big gains. Its name comes from a metaphor of slicing small slivers of salami sausage to create enough pieces to create a sandwich.

In one of the classic examples of a salami attack, an attacker shaves a penny off each of many financial transactions. Any individual who loses a penny is unlikely to notice or even complain about the loss. However, when attackers use the same tactic thousands or millions of times, they can pilfer a substantial amount of money. For example, programmers in the movie *Office Space* realized that they could write a program to shave fractions of pennies from transactions where they work. Unfortunately, they misplaced a decimal and ended up taking too much money, resulting in some comical results.

Salami attacks can also be used to gain information from databases. Instead of getting the entire database at a time, attackers can get small pieces of data and then combine them to reproduce the full picture. Espionage works this way when spies try to learn as much as they can using multiple sources. Similarly, social engineers (described later in this chapter) often try to gather small pieces of data from multiple sources.



TIP Chapter 11 delves into databases a little deeper, including the concept of inference. In an inference attack, an attacker collects multiple pieces of small, seemingly inconsequential information, but is able to put this together to infer or deduce sensitive information.

Man-in-the-Middle

Man-in-the-middle (MITM) attacks perform electronic eavesdropping by placing a system between two other parties. This third system captures the data sent between the two computers. For example, if one computer is exchanging data with a server, an MITM attack can capture all the data between the two systems. The attack is transparent to the two systems.



EXAM TIP A packet sniffer such as Wireshark is used in an MITM attack. The attacker captures packets and then analyzes the data to learn information. If systems transmit data in cleartext, the attacker can capture and read it.

An active MITM attack can be quite sophisticated, and Figure 5-7 shows an example. A user is connecting to a web server over the Internet and the user's web browser indicates that a Transport Layer Security (TLS) session is established with HTTPS. However, instead of having a secure TLS session with the web browser, the user has a secure TLS session with the MITM system.

The MITM system establishes two TLS sessions. One session is secure between the user and the MITM system, and the other session is secure between the MITM system and the web server. The MITM system is able to decrypt all of the data, allowing attackers to read it.

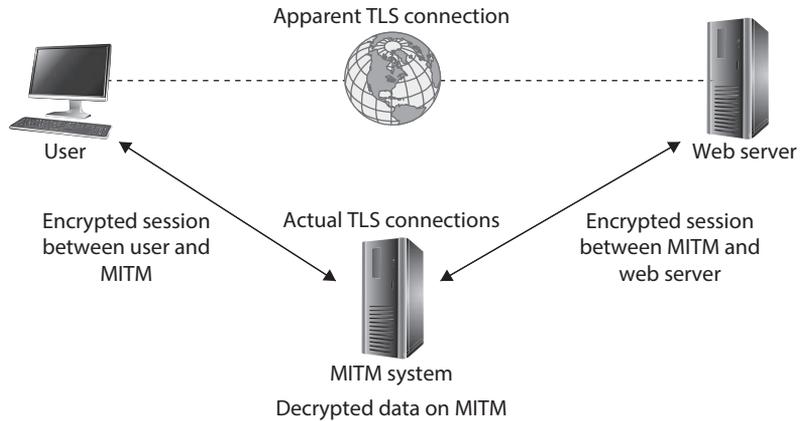


Figure 5-7 MITM attack using TLS

It's difficult to set up such an attack without control of either the user's internal network or the web server's internal network. Some organizations have set up an MITM system internally to track activity of users within the organization. For example, if the MITM system is established internally, it allows personnel to track all of the internal users' online activity. Before doing this, an organization would typically let employees know via an acceptable use policy that their online activity is tracked and that certain activity is prohibited.

If the web server is hosted at a disreputable hosting company, the hosting company may be able to implement an MITM system there. It can then capture user data such as credit card data and then run up charges on these credit cards.



TIP MITM attacks can be active or passive. In an active MITM attack, the attacker tries to capture data from specific systems. In a passive attack, the attacker attempts to capture data from any system.

Session Hijacking

Session hijacking is another sophisticated attack where an attacker captures information from an ongoing TCP/IP session and attempts to take over the session by impersonating one of the parties. Such an attack is possible with HTTP sessions where cookies are used. If successful, the attacker may gain access to either of the systems in the original session or launch a DoS attack on one of the systems to take it out of service.

The session hijacking attack often starts with a sniffing attack where the attacker gathers information about an ongoing session. The gathered information includes cookies passed back and forth between the web client and the web server.



TIP A cookie is a text file stored on a user's system. When a user visits a website, the website is able to read the cookie and use it to identify the user. The cookie also includes a session ID used to identify the current session to the web server.

As an example, Firesheep is a Firefox extension that allows attackers to listen for and exploit session cookies transmitted over a wireless network. An attacker can sit in a public wireless hotspot, capture session cookies, and use them to take over sessions such as Facebook and Twitter sessions. Faceniff is a similar application that runs on Android systems.

Replay

A replay attack starts as an eavesdropping attack where the attacker uses a sniffer to capture the data. However, the replay attack goes a step further and uses the information to impersonate one of the systems in the session.

For example, two computers may establish a session and exchange authentication information. An attacker will capture this information and use it later to impersonate one of the systems.

Several methods help prevent the success of this type of attack. For example, Kerberos (a network authentication protocol) uses time stamps and requires systems to be within five minutes of each other. These time stamps restrict the time that an attacker can replay captured data. Additionally, many systems use unique session tokens within secure sessions to thwart replay attacks. For example, Challenge-Handshake Authentication Protocol (CHAP) uses a nonce (a number used once) to create session tokens. Because CHAP doesn't reuse these session tokens, attackers cannot use them in a replay attack.

Smurf and Fraggle Attacks

A smurf attack broadcasts ICMP ping packets to multiple computers on a network but spoofs the source address using the IP address of the attacked system. An ICMP packet normally includes the IP address of the sender in the source IP address field. However, by replacing the IP address with the victim's IP address, the ICMP packet appears to come from the victim's computer. By broadcasting the ping, all systems on the subnet receive the echo and respond by flooding the attacked system with echo replies.

Attackers often use an amplifying network in a smurf attack by sending a directed broadcast ping to the amplifying network. A directed broadcast ping goes through a router to the target network and then broadcasts the ping to all the computers on the target network. Each computer on this network then sends pings to the victim's computer. However, most routers block directed broadcasts, protecting a network from being used as an amplifying network.

A fraggle attack is similar to the smurf attack, but instead of using ICMP packets, it uses User Datagram Protocol (UDP) packets for the attack. It sends packets to UDP port 7 or UDP port 19. Port 7 is the echo port, which works similarly to a ping. Port 19 is the Character Generator Protocol (chargen) port. When chargen is enabled on a system, it will respond

with a random character each time it receives any traffic on the port. However, chargen is rarely enabled on current systems.

Software Security as a Countermeasure

Many organizations host their own websites, and developers can take steps to protect these websites from various attacks. Successful attacks can infect the websites and infect user systems that visit the websites. Software security includes steps that developers within an organization can take to ensure that their software applications (including web applications) are secure.

Years ago, it wasn't common for an organization to develop its own software applications. However, now it is common for an organization to create simple web applications on its own instead of hiring an outsider. Many organizations let internal employees create them, and this practice is often successful. However, if the website developer isn't aware of basic security measures, including input validation, application reviews, code signing, and sandboxing, security problems can arise.

Using Input Validation as a Countermeasure

Input validation checks data before using it within the application. The goal is to verify that the data is valid and ensure that an application does not use invalid data. This prevents potential problems within the application and helps prevent attacks.

As an example of input validation, imagine an application needs a user to enter a number between 1 and 100 that the application will use in a mathematical equation. Input validation does a check to verify that the user actually entered a number in this range. If the user enters 0, a number over 100, or any non-numerical characters, it fails the input validation check, and the application won't use the entered data.

Applications should use input validation for any type of data. If users need to enter their first name, the application can check to ensure that the entered name is less than 35 characters and includes only valid characters for a name. Similarly, if users need to enter a five-digit ZIP code, input validation verifies that the input includes five digits.

Input validation doesn't verify the accuracy of the data. Instead, it only determines if it is valid. As an example, if you can enter 23456 as your ZIP code, a web application using input validation will accept it because it is a valid ZIP code. However, the application doesn't know if this is an accurate ZIP code for you. More advanced accuracy checks can verify that the ZIP code you entered matches the city and state you entered, but these are accuracy checks instead of validity checks.

If you've ever filled out a form on a web page, you've probably seen input validation in use. After you enter the data and submit it to the website, the web page checks your input. If something is wrong, the site redisplay the web page, often using something like a red asterisk (*) to show what items need to be addressed.

Input validation is valuable in preventing different types of attacks, such as buffer overflow, SQL injection, and cross-site scripting attacks. These attacks are discussed later in this chapter.



EXAM TIP Applications use input validation to check user input before using the data. Input validation often checks for specific characters such as < and > to prevent cross-site scripting attacks and can help prevent buffer overflow attacks.

Client Side vs. Server Side Input Validation

Input validation can be done on the client side and/or on the server side. With client-side input validation, the web page includes embedded code that validates the user input. If the user enters incorrect data, the code within the web page validates it and modifies the display to let the user know the data isn't valid.

Client-side validation saves a round trip to the server over the Internet. In other words, the web page doesn't send the input over the Internet to the web server, requiring the web server to send a page back to the user to correct the input. However, a malicious attacker can modify the web page on the client side and bypass the client-side validation. Because of this, client-side validation by itself doesn't provide reliable protection.

Server-side input validation validates the data when the server receives it. If the input is invalid, the server rebuilds the web page and sends it back to the client. Web pages commonly use both client-side input validation and server-side input validation. The client-side input validation saves bandwidth and processing power on the server. The server-side input validation provides a final security check before using the inputted data.

Application Review as a Countermeasure

An application review attempts to discover vulnerabilities with an application before an organization releases it. One of the basic elements of a review is testing. A tester verifies that the application works as it should, and tests for any bugs or vulnerabilities.

Application reviews also include a review of the source code. This is a line-by-line examination of the source code of an application. The primary goal is to discover potential issues that developers can correct before the organization releases the application. Some organizations also use application reviews as a peer review to help developers improve their skills.

Code Signing as a Countermeasure

Code signing, or digitally signing software, provides security for the organization hosting a website by associating a certificate with the software. It also provides users with assurances of who wrote the software and that the software has not lost integrity. If malware infects digitally signed software, it will fail a check and alert the user. In some cases, the software is within a web page and the browser will raise an alert. In other cases, the operating systems will detect the failed check during the installation and raise the alert. Note that code signing verifies both the identity of the author and file integrity.

Vendors that release software purchase a certificate from a public certificate authority (CA). This certificate is associated with the software file and the vendor that released it. It includes the name of the vendor and a hash of the file.

When users download the software file, they also download the certificate. When they install the software, the system verifies the certificate matches the file and that the cer-

tificate is valid. Chapter 14 discusses how a system validates a certificate, but in short, it queries the CA to ensure the CA issued the certificate, and that the CA has not revoked it. The system then displays a dialog box that displays information about the certificate. If the code doesn't have a certificate, the dialog box indicates the vendor is unknown. If malware infects the file, it changes the hash, and the certificate will no longer recognize the file. Instead, the certificate reports an error.

Although many legitimate software vendors are willing to purchase certificates to provide assurances to their customers, malware authors are not. If criminals purchased certificates for malware they wrote, law enforcement would be able to identify the criminals relatively easily. However, criminals don't want to lead law enforcement personnel to their door, so they don't purchase certificates and historically haven't digitally signed their malware.

However, in the past few years, criminals have been signing malware with code signing certificates that they stole from others. Symantec introduced Extended Validation (EV) Code Signing certificates, and the CA/Browser Forum released standards for these in 2012. EV Code Signing certificates provide a more thorough background check with the intent of providing users with a higher level of confidence for code signing certificates.

With this in mind, it's a safe bet that digitally signed software using EV Code Signing certificates is relatively safe. However, other code may not be safe. This doesn't mean that all unsigned software is malicious, but unsigned software should be viewed with some suspicion. Additionally, software using standard (not EV) code signing methods should be viewed with some suspicion.

Both Java applets and ActiveX controls (commonly used in websites) support code signing. Most browsers have settings that can restrict the use of unsigned software, or at least provide users with a warning if a website tries to run unsigned software.



TIP Configuring web browsers to block unsigned software greatly increases security. Unfortunately, some legitimate websites that use Java applets and ActiveX controls don't go through the expense of purchasing certificates for them. In other words, restricting a web browser to only digitally signed software increases security but often reduces usability.

Malicious ActiveX controls can gain a significant amount of control over a Windows operating system. ActiveX controls are largely limited to Microsoft Windows systems, but some ActiveX controls can run in Internet Explorer running on other operating systems. Java applets are theoretically portable to any operating system. However, due to differences in operating systems, the portability isn't 100 percent.

Sandboxing as a Countermeasure

Sandboxing runs computer programs in isolated areas of memory as a security control. Antivirus programs use sandboxing to check new or untrusted applications. If antivirus software determines that a program is malicious, it blocks the application from running normally. This prevents a malicious program from infecting a computer.

Security professionals often use virtualization to test applications. For example, they create a virtual machine, and then isolate it from the physical computer and other virtual machines. They can then run the application within this isolated virtual machine. If the application causes problems, the problems are limited to the virtual machine.

Buffer Overflow Attacks

A buffer overflow occurs when an application receives more data than it expects and is unable to handle it gracefully. Applications have areas of memory called buffers that they use for temporary data storage. These buffers are specifically defined areas of memory, and applications are normally unable to access memory spaces beyond the buffers. However, there are times when an application gets overwhelmed by the amount of data sent to it, allowing it to access memory beyond the defined buffer.

A buffer overflow error exposes normally protected memory areas. If attackers can exploit a buffer overflow error, they are able to write additional data into the overflow area. Attackers often include a large volume of no operation (NOOP) commands in an attack to position the memory pointer. Attackers can then add malicious code after the NOOPs, allowing them to insert the malicious code into memory.



EXAM TIP A buffer overflow problem starts as a programming error. However, if attackers detect it, they can launch a buffer overflow attack. Successful buffer overflow attacks allow attackers to gain elevated privileges and insert malicious code into normally protected memory areas.

As an example, an application may expect a user to enter some specific data, such as a number between 1 and 10. An attacker may try to send different data, such as an escape sequence and multiple NOOP commands, to see what happens. If the different data is not expected, it can cause an error. If the application doesn't handle the error gracefully, the attacker can try to repeat the error but add more data. This additional data can include a malicious script that installs itself on the server.

Proper error handling within an application ensures that it can handle errors gracefully. Error handling routines provide user-friendly errors to users and log details giving administrators insight into what happened. They also ensure the application doesn't crash.

Input validation helps prevent buffer overflow attacks. Additionally, keeping systems updated protects against buffer overflow attacks. When a buffer overflow vulnerability occurs, vendors release patches that fix the problem. Patched systems are protected against known buffer overflow attacks. Of course, if patches are not applied, the system remains vulnerable.



EXAM TIP A key to keeping a system secure is to keep it updated with current patches and fixes. This protects against buffer overflow attacks and many other types of attacks.

Injection Attacks

Injection attacks include SQL injection and command injection attacks. In either type of attack, the attacker attempts to inject code into a web application. If successful, the code runs a command outside of the web application. Input validation is a primary prevention against these types of attacks.

SQL Injection

SQL (Structured Query Language) is the language used to communicate with databases. SQL statements provide the ability to read, write, modify, and delete data within a database. In many cases, programmers create richer websites by interacting with a database, and they use SQL statements to do so.

Normally, this works fine. However, many malicious attackers know about SQL and they can sometimes use these same SQL statements to interact with a database through the website. As an example, when entering data into a website form, an attacker can enter a specially crafted string of code that includes a valid SQL statement. This allows attackers to access the backend database. To get a better idea of how a SQL injection attack can access data in a database, it's useful to understand how dynamic SQL works.

Dynamic SQL Imagine a website that expects a user to enter a username and password into two text boxes on a website. If the credentials match, the website can retrieve the user's name from a table named Authentication. The two text boxes on the web page are named txtName and txtPassword, and the following dynamic SQL statement will retrieve the user's first name:

```
SELECT FirstName FROM Authentication
WHERE login = 'txtName.Text' and password = 'txtPassword.Text';
```

The statement will take the data entered in both text boxes and dynamically add it to the SQL statement. For example, if Darril enters **Darril** as the username and **P@ssw0rd** as the password, the SQL statement is interpreted like this:

```
SELECT FirstName FROM Authentication
WHERE username = 'Darril' and password = 'P@ssw0rd';
```



TIP This example shows how dynamic SQL can be used. Although such code will work, and developers who don't know any better sometimes use it, it is not recommended. Attackers can exploit sites using dynamic SQL statements with a SQL injection attack.

SQL Injection Attack An attacker with a little knowledge of SQL can enter specific strings of characters to access the database, and possibly cause damage. Two pieces of knowledge about SQL syntax are important here:

- A semicolon (;) indicates the end of a SQL statement.
- Two hyphens (--) indicate a comment that is ignored.

With this knowledge, the attacker can enter different data in the txtName and txtPassword text boxes discussed in the previous section. Suppose the attacker enters **whoever' OR 5=5; --** in the txtName text box and **password** in the txtPassword text box as part of a SQL injection attack. Dynamic SQL uses these entries to create the following statement:

```
SELECT FirstName FROM Authentication
WHERE username = 'whoever' OR 5=5; -- and password = 'password';
```

The WHERE clause checks for a login name of *whoever* in the Authentication table, but also includes an OR statement checking to see whether 5 equals 5. Because 5 always equals 5, the WHERE clause is always true. The statement retrieves the first name of all the users in the Authentication table. The semicolon ends the statement, and the two hyphens mark the rest of the text as a comment, which SQL ignores.

Although retrieving first names from the database isn't necessarily a problem, a SQL injection attack allows an attacker to inject any SQL statements. Attackers can retrieve all the data from the database, modify data, and even delete data.

Input validation helps prevent this type of attack. For example, a username should not include double hyphens or semicolons, and an input validation check can reject input using these characters.

Database developers also use stored procedures instead of dynamic SQL statements. A *stored procedure* is a group of SQL statements that executes as a small program. Instead of plugging the user input directly into the SQL statement, the data is sent to the stored procedure. The stored procedure handles the data differently and prevents a SQL injection attack.



EXAM TIP Input validation and stored procedures help prevent SQL injection attacks.

Some SQL injection attacks have infected systems with malware. For example, the LizaMoon SQL injection attack infected websites in 2010 and 2011 with scareware. When users visited the infected websites, they saw a pop-up window that indicated their systems were infected with malware. It encouraged users to install rogue antivirus software named Windows Stability Center (which is not a Microsoft product). Users who installed the bogus antivirus software were then encouraged to purchase the full version to remove reported malware.

Command Injection

In a command injection attack, an attacker attempts to insert operating system commands into web page forms. If successful, the command executes on the system hosting the web page.

A simple operating system command is **dir**, which lists the files in an operating system directory or folder. An attacker may be able to insert a specially crafted string of characters into a web page form, followed by the **dir** command. The string of characters provides access to the operating system command shell and then executes the **dir** command.

Of course, the attack can be much more malicious than just retrieving a list of files. If successful, the attacker is able to execute any command supported by the operating system. As with other injection attacks, input validation is the primary prevention.

Cross-Site Scripting

Cross-site scripting (XSS) is a relatively common attack that attempts to inject HTML or JavaScript code into a web page. The attacker modifies the code on the website, and when the user goes to the website, the code executes on the user's system.



TIP XSS is used instead of CSS as an acronym for cross-site scripting because CSS is the acronym for Cascading Style Sheets. Web pages use Cascading Style Sheets to control the look and formatting of HTML data.

In most cases, the website is a legitimate one, but after the attack, it can become a malicious website for visiting users. The attack doesn't cause direct problems for the website, but it can cause problems for the users and indirect problems for the website.

XSS attacks attempt to steal sensitive information from users' systems, such as cookies or other information. In some cases, the cookies can include session information, and an attacker can use this to launch a session hijacking attempt. If the website developer used the cookie to store user information such as passwords or credit card data, the attacker now has access to this information.

As an example, attackers once discovered an XSS vulnerability on Twitter.com. The attackers used XSS to cause pop-up windows to appear when the user's cursor hovered over a link. In many cases, the pop-up windows displayed advertisements and links for hardcore porn sites. This provides a good example of how an XSS attack can cause indirect problems for an organization. This attack didn't affect the functionality of Twitter. However, it was embarrassing and drew quite a bit of negative publicity even though Twitter corrected the problem soon after employees learned of it.

As with other types of injection attacks, the best protection is server-side input validation. In this case, the website application can filter out any attempts to include HTML or JavaScript tags using the < and > characters.

Cross-Site Request Forgery

Cross-site request forgery (CSRF, pronounced as *sea surf*, and sometimes identified as XSRF) includes "cross-site" in the name, but it is different from an XSS attack. An XSS attack attempts to download malicious code from the website to the user's system. A CSRF attack attempts to send a malicious command from the user's system to the website. Both occur without the user's knowledge.

In a successful CSRF attack, the attacker tricks the user into clicking a malicious link that includes a command. The website treats the command as if the user sent it and executes it.

As an example, consider an HTML query. HTML queries can include additional data. For example, if you launched Google, you could enter the following text into Google to

do a search: **Darril “Security+”**. You can also do the same search with the following query embedded in a hyperlink:

```
https://www.google.com/search?as_q=Darril&as_epq=%22Security%2B%22
```

This query includes several elements worth pointing out:

- **?as_q** The question mark character indicates that the appended HTML link includes a query, and as_q indicates the query is using Google’s advanced search function.
- **=Darril** This equal sign identifies a search term, which is Darril in this example.
- **&as_epq=%22Security%2B%22** The ampersand indicates another search term, and as_epq indicates it is using Google’s advanced search function with quotes. It includes %22, which is the hexadecimal code for a quote symbol, and %2B, which is the hexadecimal code for the plus symbol. It looks like “Security+” in the query.

Websites often use the same type of HTML queries to send data entered in one page to a second page. For example, if mcgraw-hill.com allowed users to provide information such as their name and e-mail addresses, it could also have a web page that allows users to change their e-mail address. Users could log on, enter their new e-mail address in the form, and click a button to submit the change. When users click the button, it sends a query similar to this:

```
http://mcgraw-hill.com/change?action=set&key=e-mail&value=newe-mail@gmail.com
```



NOTE This is only an example and will not work on the McGraw-Hill website.

In some cases, the user doesn’t even need to be logged on to take this action. Many websites use cookies. When the user revisits a website, it reads the cookie and logs the user back on. If a user clicks a link or enters the string into his or her web browser, the website can identify the user by reading the cookie and taking the requested action.

An attacker can use this information to craft a different HTTP link to change the user’s e-mail address to the address of the attacker. Assuming the previous McGraw-Hill query actually worked, the attacker’s query would look something like this:

```
http://mcgraw-hill.com/change?action=set&key=e-mail&value=attacker@attacker.com
```

Once attackers craft the HTTP link, they only need to encourage the user to click it. The link can be embedded in a phishing e-mail or in any type of HTML object, such as an image that includes a link. When the user clicks the link, it sends the command to the website.

Changing an e-mail address may sound trivial. However, if the site doesn’t use security to validate the user before making a change, attackers may be able to change other types

of information, such as the user's password. In some cases, CSRF attacks may enable attackers to make purchases on behalf of the user.

Websites protect against these attacks by verifying the user before making such a change. Most professional web developers know better than to take some actions without validating the user again. In other words, they use a separate step to log a user on instead of trusting the cookie actions. As an example, if you've ever bought anything from Amazon.com using your computer, your computer has an Amazon cookie that allows Amazon to identify you when you visit. However, if you try to make a purchase or modify any of your account details, Amazon challenges you to enter your credentials again. This prevents a successful CSRF attack.

One of the best protections users can take is to avoid clicking unknown links. This is especially true for any links received via e-mail from an unknown source.



EXAM TIP An XSS attack exploits the user's trust of a website and downloads code onto the user's system. A CSRF attack exploits the trust that a site has in the user's browser and attempts to send unauthorized commands from the user's system to the website.

Password Attacks

Attackers often try to learn passwords through different methods. If the attacker is able to learn the username and password of an account, the attacker can then impersonate the user by logging on with the user's credentials. A password cracker is a program that tries to discover, or crack, passwords.



EXAM TIP Strong, complex passwords can thwart most password attacks. A strong, complex password is sufficiently long and includes multiple character types such as uppercase letters, lowercase letters, numbers, and special characters. Chapter 2 covers passwords in more depth.

The following list describes some common password attack methods:

- **Brute-force** In a brute-force attack, the attacker tries all possibilities until getting a match. This can be quite time consuming when performed manually, but password-cracking tools automate the attempts. Brute-force efforts generally take the most time and effort when compared with other password attacks. This attack can be thwarted by using strong passwords of sufficient length and with multiple character types. The longer and more complex a password is, the harder it is to crack with a brute-force attempt.
- **Dictionary** A dictionary attack attempts to use all the words in a dictionary or a database of different words. Password dictionaries usually include words from multiple languages in addition to commonly used passwords such as 123456, iloveyou, qwerty (the first five letters next to the TAB key), 7777777, letmein (let me in), trustno1 (trust no one), and qazwsx (the first six letters from top

to bottom on the left side of the keyboard). Strong, complex passwords thwart dictionary attacks.

- **Social engineering** The simplest way to get a password is simply to ask for it. Although most users won't give up their password easily, social engineers often trick users. For example, phishing attacks (described later in this chapter) often ask users for their username and password, and some users respond.



TIP Dictionaries are available in multiple languages, so using words from different languages doesn't make a password more secure.

A rainbow table is a more sophisticated attack used to learn passwords. It uses comparative analysis techniques to identify passwords. To make this attack easier to understand, it's worthwhile examining each of the steps in a comparative analysis.

One of the ways that passwords are protected is by hashing them. A hash is simply a number created by executing a hashing algorithm against the password, and the hash will always be the same if the password is the same. For example, imagine a password is P@ssw0rd and the hash of the password is 1A2B3C. No matter how many times you execute the hashing algorithm against the password, it will always create a hash of 1A2B3C. (The actual hash would be more complex, but I've shortened it here for brevity.)

Now imagine that an attacker has captured a hashed password either by hacking into a password database or by capturing packets with a sniffer. In either case, the attacker now knows that the hash of a password is 1A2B3C. The attacker can use the following steps to discover the password from the hash:

1. The attacker guesses a password.
2. The attacker calculates the hash of the guessed password.
3. The attacker compares the calculated hash with the captured hash.
 - a. If the hashes are the same, the attacker knows the guessed password is the actual password.
 - b. If the hashes are not the same, the attacker repeats the steps.

The most time-consuming part of this process is guessing another password and then calculating the hash. However, a rainbow table shortens this time considerably. A rainbow table includes two columns. One column is a list of guessed passwords and the matching column includes the hashes of the guessed passwords. A program compares the captured hash with each of the hashes in the rainbow table. When it finds a match, it retrieves the actual password from the second column.

Attackers create rainbow tables using programs such as RainbowCrack and character set definition files that define the keyspace for the password. A typical character set definition file includes the following character set definition lines:

- **numeric:**
[0123456789]

- **alpha:**
[ABCDEFGHIJKLMNOPQRSTUVWXYZ]
- **alpha-numeric:**
[ABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789]
- **loweralpha:**
[abcdefghijklmnopqrstuvwxyz]
- **loweralpha-numeric:**
[abcdefghijklmnopqrstuvwxyz0123456789]
- **mixalpha:**
[abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ]
- **mixalpha-numeric:**
[abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789]
- **keyboardascii-32-95:**
[!"#\$%&'()*+,-./0123456789;:<=>?@ABCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz{|}~]
- **ascii-32-65-123-4:**
[!"#\$%&'()*+,-./0123456789;:<=>?@ABCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`{|}~]
- **alpha-numeric-symbol32-space:**
[ABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789!@#%^&*()-_+=~`[]|\:;'"<>.,?/]

The rainbow table generation program then creates passwords using all of the characters in each character set. For example, the attacker can specify a password length up to 14 characters. The program will then create all possible combinations using between 1 and 14 characters, and using only the characters in the numeric data set (0123456789). As it creates the password, it also creates the hash and stores both in the rainbow table. When finished, the program does the same with each of the remaining character sets.

It can take several hours to create a rainbow table. The actual time depends on how many character sets are used, the length of the passwords, the hashing algorithm, and the processing power of the computer creating them. Once the rainbow tables have been created, attackers can use them to crack passwords using a simple rainbow table lookup process.

Several rainbow tables are available using different hashing algorithms, different character sets, and targeted against specific operating systems. Attackers can download the rainbow tables from sites that have already created them, or create the rainbow tables on their own. Although rainbow tables can be very large (as big as 64GB), they significantly reduce the time required to crack a password.

A primary protection against rainbow table attacks is key stretching. Key-stretching techniques salt the password before hashing it. A salt adds a random number of bits to the hashed password. In other words, an unsalted hash of P@ssw0rd may be 1A2B3C,

but when it's salted, it could be 3602FD or just about any other hexadecimal number. Chapter 14 discusses key-stretching techniques such as Password-Based Key Derivation Function 2 (PBKDF2) and bcrypt.



EXAM TIP Some password crackers use rainbow tables to speed up the process of cracking a password. Salting hashed passwords thwarts rainbow table attacks.

Spam

Spam is unsolicited commercial e-mail (UCE) and other unwanted electronic messages sent in bulk. It is most often associated with e-mail, but attackers send spam via other means, such as instant messaging. Companies have used spam to send unwanted advertisements. However, the bigger risk is that attackers now regularly use spam to launch different types of attacks.

The cost of sending spam is low for spammers. A single e-mail can include multiple e-mail addresses, making it relatively easy to send an e-mail to thousands of recipients at a time.

Cyber criminals often rent access to botnets to send spam. For example, criminals can rent access to the Cutwail botnet *spam as a service* infrastructure, which is an interesting twist on cloud-based services. This botnet includes hundreds of thousands of zombies and automatically generates variations in the spam to evade detection by spam filters.

Criminals harvest e-mail addresses using a variety of methods. For example, some programs called *web spiders* crawl through the Web looking for the @ character to identify e-mail addresses. If you post messages on a website with your e-mail address, these spiders will likely find it and add your e-mail address to a spam list. These criminals then sell or rent the lists to others wishing to send spam.

Phishing Attacks

A phishing attack is spam with malicious intent. It encourages a user to respond with personal information or to click a link. Attackers send massive numbers of e-mails with the goal of fishing for victims. Phishing e-mails often claim to be from a legitimate company and encourage the user to either click a link or provide a response.

Here's an example of a phishing e-mail that attempts to get users to respond with their logon credentials for a bank account.

Warning! This is a security alert from YourBank. We have noticed suspicious activity on your account. To protect you, we are putting a hold on your account until we can verify certain activities. Please click the following link and validate your account. If you do not validate your account, funds in your account will be frozen and remain inaccessible.

If you click the link, it takes you to a bogus website that looks like that of YourBank, but is actually a malicious website. It includes logon text boxes where you can enter your

username and password. Of course, if you enter your credentials, the attacker will use them to remove your money.

Attackers quickly funnel money drained from accounts this way to offshore accounts. If the activity isn't discovered within 24 hours, the chances of recovering the money are extremely low. Thankfully, banks typically provide fraud insurance for personal bank accounts, so the account holders receive their money back if it is stolen. However, business accounts don't have the same protection. If an employee accidentally gives up credentials for the organization's bank accounts, attackers can drain the accounts. This can cause a financial catastrophe if the organization doesn't have fraud insurance.



TIP Phishing attempts have been circulating for years and most IT professionals, and even a lot of end users, know about them. Yet, they continue. Why? The short answer is that they continue to work. Many end users simply don't understand the risks and continue to respond to these e-mails. When users are educated and stop answering the phishing e-mails, attackers will stop trying these methods (but they're sure to try different ones). In the meantime, it's important to educate users about the risks.

You've probably received some phishing e-mails that look like they are from a bank or a company where you don't have an account. Attackers send out mass quantities of these e-mails knowing that not every recipient has an account with the bank, but also knowing that many of them do. Imagine that an attacker sends out an e-mail to 10,000 recipients but only 1,000 of the recipients (10 percent) actually have an account with that bank. If an attack reaches 1,000 potential targets, it has some great possibilities of success.

Even though phishing has been going on for a long time, it still fools many people. Google released the "Handcrafted Fraud and Extortion: Manual Account Hijacking in the Wild" study in late 2014. The study found that phishing attacks succeed between 3 percent and 45 percent of the time, with sophisticated attacks enjoying the most success. Additionally, the study found that attackers used the stolen credentials within 30 minutes after users provided them.

A phishing e-mail usually has several basic components:

- **Impersonation** A phishing e-mail impersonates a legitimate company. It may include actual graphics from the company and use similar fonts. For example, an e-mail that claims it is from PayPal may include PayPal graphics.
- **Identification of a problem** The e-mail usually indicates some type of problem that is of a concern to the reader. The phrase "suspicious activity" or something similar is often included in these e-mails.
- **A sense of urgency** Recipients are encouraged to take immediate action. Although the phishing e-mail doesn't state it this way, the goal is to get the user to click without any more thought.
- **Dire consequences** The phishing e-mail warns of some dire consequences, such as closing an account, locking out an account, or freezing funds, if the

recipient doesn't respond. For individuals doing business through a company such as eBay or PayPal, the threat of their account being closed, or locked, sounds dire. This often propels them into taking immediate action.

Phishing and Drive-by Downloads

Some phishing e-mails just want users to click a link. If they click, the link takes them to a malicious website that attempts to start a drive-by download. Notice that in this case, it doesn't matter whether a user enters any personal information or not because the attacker really only wants the user to click the link. Drive-by downloads attempt to install malware and, when successful, often join the computer to a botnet.

Some phishing e-mails mention a failed delivery attempt from a service such as FedEx, and encourage the user to click the link for more information. Other e-mails have impersonated funeral homes. They mention a pending funeral and encourage the user to click for more details. These are just two examples, but attackers have used many more methods.

Spear Phishing and Whaling

Spear phishing and whaling are variants on the phishing theme. Phishing attacks send out massive amounts of spam e-mail without knowing the targets. In contrast, spear phishing and whaling attempt to identify the targets.

A spear phishing attack targets a specific organization or group. For example, a spear phishing attack could target a military organization by sending e-mails to recipients with the appearance that Sergeant Jones sent it from within the organization. Because it looks like it's from someone within the organization, recipients are more likely to open it. The contents could state something about a missed deadline and ask the user to provide some information in an attached document. However, the attachment includes malware and attempts to install itself on the user's computer if the user opens it.

The From address provides a level of trust because it looks like it's coming from a fellow employee. Because of this, some users are more willing to open the e-mail and the attachment. In contrast, these users might not respond to a typical phishing attempt.



TIP Spear phishing can target any organization. If attackers are interested in money, they may try to target a bank. If they are interested in government espionage, they may try to target government employees. The key is that the spear phishing e-mails are targeted to a specific group.

In a whaling attack, the attacker attempts to identify an executive such as a chief executive officer (CEO), president, vice president, or manager. The attacker may be able to get the name and e-mail address information from public sources or through a social engineering attack with another employee. The subject line of a whaling attack may include words such as "Past Due" or "Lawsuit" that may get immediate attention from the executive.

If an attacker can get a response from a high-profile figure (a whale, if you will), the potential payoff is huge. For example, if an attacker is able to convince an executive to

install keyloggers on the organization's systems, the attacker can steal highly valuable information and perhaps even account information to empty an organization's bank accounts.

Vishing

Vishing is a slight variation of a phishing attack. Instead of using e-mail, the attacker uses the telephone system or possibly a Voice over Internet Protocol (VoIP) system. The goal is similar to phishing. It attempts to trick the user into providing information that the attacker can use for personal gain.

A vishing attack starts by making a call. Some use a recording to alert people of some type of fraudulent activity with their bank or credit card, and encourage them to call a phone number immediately. If individuals return the call, an automated system often prompts them to enter information in the phone's keypad, such as a credit card number, expiration date, and more.

I once received a voicemail message mentioning suspicious activity on a credit card. I called back, and the woman who answered started asking personal questions that raised alarms for me. When I expressed distrust, the woman said, "You called us," as if that was supposed to tell me everything was OK. It wasn't and I ended the call. There weren't any problems with my credit card then, but if I kept answering questions, there would have been.

In contrast, after a business trip, one of my credit cards had some actual suspicious activity. My bank called and left a message asking me to call back. When I called back, the person provided information that made it clear she was an employee of my bank before asking me to provide information about my account. In other words, we provided mutual authentication.

Because vishing uses VoIP, it often takes advantage of additional features. For example, attackers can trick caller ID so that phone calls appear to be coming from a bank or someone else other than the attacker's phone number.

Smishing

Smishing is another variation of a phishing attack, but it uses text messages such as Short Message Service (SMS) messages commonly used with many smartphones. Attackers send a text message to users, indicating that there has been suspicious activity on an account. It then prompts the user to call a phone number to resolve the problem.

When the user calls the number, he or she often hears a voice recording that prompts the user to enter information such as an account number and personal identification number. It may prompt the user to give additional information, such as their birthdate, Social Security information, and more.

Zero Day Exploits

Zero day exploits are attacks that take advantage of vulnerabilities that are unknown to and/or unpublished by the vendor. However, even though the vendor might not know

about the vulnerabilities, some attackers know about them. The two common scenarios for zero day vulnerabilities are

- **Unpublished and unknown by the vendor, but known by attackers** For example, an attacker may discover a buffer overflow vulnerability and launch an attack using this vulnerability. The attacker may share the knowledge with others, and each of these attackers can launch zero day exploit attacks.
- **Known by the vendor, but not published** At some point, the vendor learns about the vulnerability. However, it takes time to develop and release a patch to fix it. For example, an antivirus company could discover an issue with a Microsoft operating system and inform Microsoft. Microsoft will investigate and write a patch to fix it, but its personnel also need to test the patch before releasing it. Attacks that attempt to exploit the vulnerability before Microsoft publishes the patch are zero day exploits.

The danger from zero day exploits also helps to emphasize the importance of a defense in depth strategy. If the vulnerability is unknown to the public and the vendor hasn't released a patch, IT personnel won't know about it. However, other security practices can help prevent a successful attack.

For example, one of the basic elements of hardening a system, or making it more secure from the default configuration, is to disable unneeded protocols and services. Imagine that an attacker discovers a zero day vulnerability with FTP and launches an

Unpatched Systems and Zero Day Exploits

In some cases, people refer to a zero day exploit as an attack on a system shortly after the vendor releases the patch but before IT personnel apply the patch.

For example, Microsoft releases patches on "Patch Tuesday," the second Tuesday of every month. However, IT personnel don't always apply the patches immediately because they need to test the patches. They may not apply the patches until Wednesday or Thursday, or even later.

Attackers often download these patches as soon as they are available and get as much information on them as they can find. They reverse engineer the patch to discover the vulnerability the patch is fixing. Once they discover the vulnerability, they develop an attack method to exploit it.

Exploits of the vulnerability between the time when the vendor releases the patch and when IT personnel patch the system is sometimes called a zero day exploit. In this case, attacks on Tuesday or Wednesday after the patch is released may be called a zero day exploit.

Not everyone agrees with this definition though. Many security professionals say that if a patch is released, exploits against the vulnerability are simply attacks against unpatched systems. Still, this helps emphasize the importance of keeping systems patched and up to date. Many successful attacks take advantage of unpatched systems, and IT personnel can prevent these attacks by keeping systems up to date.

attack on a server. If the server has FTP enabled by default, even though it does not need FTP, it is vulnerable to this attack. On the other hand, if the server was hardened and FTP was disabled, it is not vulnerable to the attack.

Covert Channel

A *covert* action is a secret one that is not open or acknowledged, such as a clandestine spy operation, and a *channel* is a communications path between two computers. Combined, a *covert channel* is a mechanism used to transmit and receive information surreptitiously between two computers. Covert channels often use nontraditional communication paths, which helps them bypass traditional controls such as firewalls and intrusion detection systems.



TIP Covert channels can be used as a secret method to pass information between two entities for malicious purposes.

For example, TCP uses a TCP handshake process to establish a TCP session. The primary pieces of information used in these packets, besides the source and destination data, are the SYN and ACK flags. However, there is room elsewhere in the TCP packet, such as in the protocol ID field, to add more data. For example, one person could add ASCII codes for individual characters in the protocol ID field before sending the TCP packet to another person. If both people know the packets include this data, they can use it as a covert channel to share information.

The ability to transfer small pieces of data in the protocol ID field is a well-known tactic. Many firewalls and intrusion detection systems are able to check packets for this additional data and identify the covert channel. However, attackers can use other covert channels that aren't known to others.

Wireless Attacks and Countermeasures

Wireless networks have become very popular due to their convenience. However, convenience often has a cost, and with wireless networks, the cost is increased risk. Attackers have multiple methods they can use to attack a wireless network, but countermeasures are effective to blocking these attacks.



NOTE Chapter 3 covers wireless technologies commonly used with wireless networks.

WPA2 as a Countermeasure

Wi-Fi Protected Access 2 (WPA2) provides strong protection against attacks. WPA2-Enterprise uses an 802.1x authentication server and is more secure than WPA2-Personal. WPA2-Enterprise requires each user to log in with an account before they can access the

wireless network. WPA2-Personal only requires users to know the SSID and a password or preshared key.

Wired Equivalent Privacy (WEP) has several significant weaknesses and should not be used. Wi-Fi Protected Access (WPA) is stronger than WEP, but WPA2 is stronger. WPA2 uses Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (CCMP). CCMP is based on the Advanced Encryption Standard (AES) and provides significantly stronger security than protocols used with WEP or WPA.

WIDS and WIPS as Countermeasures

Wireless intrusion detection systems (WIDSs) and wireless intrusion prevention systems (WIPSs) monitor wireless frequency channels and can discover many problems. They monitor the 2.4-GHz and 5-GHz wireless frequency spectrums and can detect unauthorized access points and suspicious activity. This includes rogue access points, evil twins, MAC spoofing, and traffic patterns from known attacks. When they detect an unauthorized or suspicious activity, they can send alerts to notify administrators.

Rogue Access Points

A rogue access point (AP) is an unauthorized wireless AP set up within an organization's premises. Attackers connect the rogue AP to a wired network and then use it as a sniffer to capture traffic. After installing the rogue AP, attackers can collect traffic from a distance, such as in the organization's parking lot. The attacker simply uses a laptop to connect to the rogue AP and collects the traffic.

Physical security is a primary countermeasure against rogue APs. Protecting network devices (such as routers or switches) in secure server rooms or locked wiring closets prevents an attacker from connecting the rogue AP to the wired network.

Wireless spectrum analyzers can detect rogue APs. For example, Fluke's AirMagnet WiFi Analyzer Pro allows administrators to monitor all the channels in the 2.4-GHz and 5-GHz ranges to detect activity. It will identify rogue APs and monitor the overall health of the wireless network. Two other tools used to detect rogue APs are WIPSs and WIDSs, as described in the preceding section. Fluke's analyzer is an example of a WIDS.

Evil Twins

An evil twin is a type of rogue access point that appears to be legitimate. Attackers use the evil twin to trick users into connecting, and when users connect, attackers collect the data. Attackers use evil twins in one of two ways:

- **Within an organization** When evil twins are used within an organization, attackers name the APs with the same or similar service set identifier (SSID) names. Users are more likely to connect to APs with similar SSIDs. Organizations use WIDSs and WIPSs to detect these APs.
- **In a public place** Attackers set up APs in public places such as airports, hotels, coffee shops, parks, and restaurants. Users connect to get free or cheap wireless access, but the attacker uses the AP to collect their data.

The best protection against evil twins is to avoid connecting to public wireless networks. However, some basic countermeasures provide protection. They include

- Confirm that the wireless network is legitimate before connecting to it.
- Do not enter any login credentials that an attacker can capture and use.
- Do not install software from the public wireless network.

A more advanced countermeasure is to use a commercial virtual private network (VPN). For example, CyberGhost offers both free and paid versions of its CyberGhost VPN service. Users create secure connections to the VPN service through the free public wireless network. This prevents attackers running evil twins from collecting user data.



NOTE Another way of avoiding evil twins is to skip the free wireless networks and use your own smartphone for wireless access. For example, you can configure an iPhone as a hotspot and then connect your tablet or laptop to the iPhone. Many smartphones have similar capabilities, allowing users to tether another device to the smartphone. The trade-off here is that this consumes bandwidth from the user's smartphone data plan, which isn't free.

Wardriving

Wardriving is the practice of looking for wireless networks, typically from a car. Attackers can drive around with a laptop and a wireless connection. When they find a wireless network, they attempt to access it.

The best countermeasure against war driving attacks is to use secure wireless protocols. This includes WPA2. WPA2-Enterprise uses an 802.1x authentication server and is more secure than WPA2-Personal, which only requires a preshared key or password.

WPA Cracking Attacks

Attackers use WPA cracking attacks to learn the password for a wireless network. They use a wireless sniffer to capture wireless packets and look for a WPA four-way handshake (indicating a user has logged on). The handshake includes an encrypted version of the AP's password. Once attackers capture the encrypted password, they use an offline brute-force attack to discover the password.

It might take some time before a user authenticates. To speed up the process, attackers send data onto the network, causing the user's system to lose its association with the AP. This requires users to log on again, initiating the four-way handshake.

The best countermeasure against WPA cracking attacks is to use WPA2. If WPA2 is used, a strong wireless network password helps thwart WPA cracking attacks. Passwords can be up to 63 characters long.

WPS Attacks

Wireless access points typically include Wi-Fi Protected Setup (WPS). WPS allows users to configure wireless devices by clicking a button on the access point and clicking a button on the wireless device. These are typically physical buttons, but in some cases they

are logical buttons accessed via an interface. Other times, users configure the devices by entering a personal identification number (PIN). While using the buttons is an easier method for users, WPS always supports the use of a PIN.

A WPS attack uses brute force to discover the PIN. Because WPS doesn't include a lockout feature, attackers can keep guessing until they discover the correct PIN. When they guess the correct PIN, the access point responds with the SSID and password used for the access point. Reaver is an open source tool that automates the process. It can discover an eight-digit PIN in 10 hours or less.

The best countermeasure is to disable WPS when it is not in use. You can usually disable WPS within the AP interface. Unfortunately, this isn't always reliable. Some access points still respond to Reaver queries even with WPS disabled. If your AP is still responding to Reaver WPS queries with WPS disabled, check with the AP's vendor to see if it has released an update for the device.

Understanding Social Engineering

Social engineering is the practice of getting people to give up sensitive data or to perform actions they wouldn't normally perform. A social engineer uses deception and fraud to trick or manipulate unsuspecting users.



EXAM TIP Although social engineering is a low-tech attack, it can still be a very effective method of gaining unauthorized access to an organization and the organization's IT infrastructure.

Sometimes a phone call to an employee can elicit a good deal of information. Attackers can learn the names of executives through public sources such as web pages, newsletters, and company brochures. Attackers can say something like, "I'm gathering some information for so-and-so," using the name of a high-level executive. This might be enough to get the employee to give out information.

With some basic information from a first phone call, attackers can build on that with other phone calls to other departments. Repeated a few times, attackers can build their knowledge, such as identifying how the organization formats e-mail addresses (perhaps as *firstname.lastname@company.com*) and learning the format of user accounts. Attackers may make a phone call to the IT department and ask something like, "I can't remember my password. Can you help me log in?" Such an inquiry may convince an uneducated help desk technician to change the password and give it to attackers. It's also possible for an unsuspecting employee to give up his or her own username and password, allowing attackers to log on with that employee's account.



TIP Pretexting is the practice of using gathered information to create another scenario and collect additional information. The pretext, or the invented scenario, increases the possibility that the person will give up additional information to the social engineer.

Social engineers appeal to someone's inherent desire to be helpful and liked. They exploit values such as courtesy and appeal to other people's vanity. Often they take the time to build familiarity and build trust. In some cases, they can invoke the name of an authority figure, such as an executive, to encourage a person to respond. If necessary, they can use intimidation tactics and threaten undesirable action to get what they want. They often create a sense of urgency to prevent people from analyzing the activity.

One of the primary methods to counteract social engineering is to educate users in the tactics of social engineers. Once employees recognize the value of information and common tactics used by social engineers, they are better able to identify and thwart a social engineer in action. The following sections identify some common social engineering tactics.

Tailgating

Tailgating (or piggybacking) occurs when someone passes through a controlled entry without providing credentials by following closely behind someone who has provided credentials. As an example, some organizations issue proximity badges to users. Employees swipe their badge in front of a proximity badge reader and it unlocks a door. All employees should use their own badge to open the door and gain entry. However, it's possible for one employee to open the door with a badge and other people to follow that person in without using their own badges. These other people are tailgating.

Imagine the beginning of the workday as the majority of the employees are coming into the building. It's unlikely that each employee will shut the door in the face of fellow employees. If the organization is large, employees may not know who is a fellow employee and who isn't. A social engineer could easily slip into the crowd and walk into the normally secure building without raising any suspicion.

Most people are courteous and consider slamming the door on someone else to be rude. When people have the choice of being either courteous or secure, they often choose the courteous path. Attackers know this. Often just a friendly smile is all that's needed if someone gives the social engineer a questioning look. Similarly, an attacker can load her arms with books or boxes and approach the door. A friendly employee may actually hold the door open for her to enter.

Educating users is a great start against tailgating, but it isn't always enough. Another way to stop tailgating is with a mantrap. The formal definition of a *mantrap* is a large enclosure that allows entry through one side and exit through the other, and can actually lock the person inside. Some are large steel bars formed as a revolving door that looks very similar to a cage.

However, instead of using cages to control entries, many organizations simplify this to turnstiles similar to what you see in a subway or bus station. Employees swipe their badge and the turnstile releases to allow one person through. When using a mantrap, even if it's a simple one like a turnstile, employees cannot keep the entrance open for a potential attacker, and it thwarts tailgating.



EXAM TIP Tailgating is a social engineering tactic where one person uses credentials to gain entry, while a second person follows closely behind without providing credentials. A mantrap protects against tailgating.

Impersonation

Another social engineering tactic is impersonation, where the social engineer impersonates someone. Impersonation can be over the phone, such as by invoking the name of someone in authority. It can also be in person, such as impersonating a technician.

For example, it's relatively easy to get a uniform that looks like it's from another company, such as the phone company. The social engineer can don the uniform, enter the building, and let a receptionist know he or she is there to fix some fictional problem. An uneducated employee might lead the attacker directly to a wiring closet, where the attacker can install a rogue access point.



EXAM TIP Impersonation is also known as masquerading or spoofing. In social engineering, the social engineer is attempting to masquerade as someone else.

Dumpster Diving

Many attackers gain information by sifting through someone else's trash, a practice commonly known as dumpster diving. Dumpster diving can provide significant returns depending on how much information an organization or an individual throws away.

For example, if an organization regularly throws away papers that include information such as customer names, addresses, and credit card data, an attacker can retrieve this from the trash and use it. Similarly, if an individual throws away credit card offers, blank checks from a credit card company, or other personal data, an attacker can retrieve it and use it.

As a best practice, any information that contains any type of personal information should not be thrown away. Instead, these papers should be shredded using a crosscut shredder or burned to ensure they cannot be read.

Shoulder Surfing

Shoulder surfing is the practice of looking over someone's shoulder to gain information. For example, an attacker may try to watch someone enter a username and password to learn the user's credentials, or watch someone enter the numbers for a cipher lock on a door, a PIN for a badge, or even a PIN for a debit card.

Pharming

Pharming is an attack that redirects users to bogus websites. It manipulates host name resolution methods so that host names resolve to a different website.

As an example, users may attempt to go to yourbank.com, but a pharming attack redirects them to a bogus website. The attacker creates the bogus website to look exactly like yourbank.com, and tries to dupe users into entering their credentials. Other attacks redirect users to malicious websites that attempt drive-by downloads, or attempt to trick users into installing a Trojan horse. Some attacks simply redirect the users to sites that sell a product such as pharmaceutical drugs.

Host names identify computers, and name resolution methods resolve the host names to IP addresses. For example, if you want to go to Google, you can type **google.com** in your web browser instead of Google's IP address, and the system resolves google.com to an IP address. The primary method of resolving a host name to an IP address on the Internet is Domain Name System (DNS), and a secondary method uses the hosts file.

In some cases, attackers manipulate data in DNS to redirect the users. A DNS cache poisoning attack is one way attackers have done this in the past, due to some previous vulnerabilities with DNS. DNS cache poisoning is often referred to as the first pharming attack.

In other attacks, the attacker manipulates the hosts file. Operating systems look at this file first for name resolution. If it includes a host name and IP address mapping, the system uses this IP address rather than querying DNS for name resolution. Microsoft systems store this file in the c:\windows\system32\drivers\etc folder by default, and many malicious software attacks have modified this file. In one case, malware added a bogus entry for Microsoft's Update site in the hosts file, preventing computers from downloading updates.

Social Networking Attacks

Attackers are taking advantage of the popularity of social networks such as Facebook, Twitter, and YouTube to launch social engineering attacks. Attackers craft e-mails that look exactly like they came from the social network site and send them out to users.

For example, Facebook often sends notifications to users indicating that they have messages or pending invites and includes a link to the Facebook site. The attacker uses the same template letting the user know a message is waiting, but includes a link to a malicious website. If the user clicks the link, the website may attempt a drive-by download or encourage the user to install a Trojan horse. In other cases, the site is a fake one claiming to sell cheap drugs or casino gambling access or some other site intended to get the user to enter credit card data.

User Awareness as a Countermeasure

A primary method of combating social engineering attacks is by raising user awareness through training and other methods. When users recognize the tricks that social engineers use, they are less likely to fall for them.



EXAM TIP A primary purpose of user awareness campaigns and training is to change user behaviors from unsafe actions to safe actions. For example, if users recognize the malicious nature of phishing e-mails, they are less likely to click the link or respond.

User awareness training can't just be a one-time event, either. Instead, users need occasional retraining. Attackers are constantly modifying their attacks and attack methods. Likewise, users need updates on these different tactics, in addition to reminders of risks.

The following are some of the common methods used to educate and reeducate users:

- **Acceptable use policy** Most organizations include an acceptable use policy that informs users of the organization's expectations regarding the employee's use of computing resources. It lets users know what their responsibilities are and can include basic information on the risks of social engineering attacks. Employees normally review and acknowledge an acceptable use policy when they are hired and then periodically afterward, such as annually.
- **Initial training when hired** Many organizations include initial security training as part of a new employee's indoctrination. This could be a short presentation, a short video, an online presentation, or even a written document for the new employee to read and acknowledge. As risks and threats change, the organization modifies the initial training.
- **Annual refresher training** Refresher training reminds personnel about basic risks and trains them about new and current attack methods. Organizations often require employees to read and acknowledge an acceptable use policy at the same time.
- **Newsletters** If an organization sends out regular newsletters, it's relatively easy to add a short article or a sidebar about security. It reminds users about security while providing information on current threats.
- **Periodic e-mails** Just a simple e-mail reminding users of certain risks, with an example of a recent way that a victim was tricked into giving up information, can be useful.



NOTE Many training companies specialize in providing content for IT security. Organizations can purchase access to online training, computer-based training delivered to the desktop, videos, e-mail-based messages that an organization can tailor and send out, and even articles that an organization can include in a newsletter. By purchasing outsourced materials, the company is able to get access to current, relevant content.

Chapter Review

Traditional definitions indicate that hackers aren't malicious, but crackers are. However, the media generally uses the term *hacker* to identify anyone who launches attacks on computers or computer systems with malicious intent. Some security professionals separate the good guys from the bad guys with terms such as *white hats* and *black hats* and refer to anyone who attacks with malicious intent as an attacker. An advanced persistent threat (APT) is a state-sponsored group that has both the capability and intent to carry out successful attacks.

Attackers are constantly modifying their attack strategies just as IT security professionals are constantly modifying their countermeasure strategies. Even though attacks have changed over the years, many attacks share common characteristics. By understanding

past attack methods, it is easier to respond to new threats and understand the need for certain countermeasures. Basic countermeasures include patching systems, hardening systems, and increasing user awareness.

A denial of service (DoS) attack launches an attack on a single system from a single system. A distributed DoS (DDoS) attack launches an attack on a single system from multiple systems. A botnet is a group of computers (called zombies) controlled by an attacker through a command and control center. A botnet can control tens of thousands of computers, and some have controlled over a million computers. Criminals rent out the botnets to others for money, and these botnets are used to launch attacks and send out spam.

Organizations that host their own websites need to implement software security. Countermeasures include input validation, application review, code signing, and sandboxing. Input validation is the most important countermeasure. It checks all input data before using it to ensure it is valid. It helps prevent many attacks, including buffer overflow attacks, SQL injection attacks, cross-site scripting (XSS) attacks, and cross-site request forgery (CSRF) attacks. A buffer overflow attack sends unexpected data to an application and exposes normally protected memory areas. SQL injection attacks send embedded SQL statements as input to access backend databases. An XSS attack exploits the user's trust of a website and downloads code onto the user's system. A CSRF attack exploits the trust that a site has in the user's browser and attempts to send unauthorized commands from the user's system to the website.

Phishing is a form of social engineering that uses e-mail. The attacker sends out massive amounts of e-mails hoping that someone responds by either clicking a link or responding to the e-mail with sensitive information such as a username and password. Simply by clicking the link, the responder can install malware on their system through a drive-by download. Variations on phishing attacks include spear phishing (targeted at an organization), whaling (targeted at a specific person, such as an executive), vishing (using telephones or VoIP), and smishing (using text or SMS messages).

Social engineering is the practice of using primarily nontechnical means to get people to give up sensitive data or perform actions they wouldn't normally perform. A social engineer uses deception and fraud to trick or manipulate people into giving up information they wouldn't normally give up. Tailgating occurs when one person enters a controlled space without providing credentials by following closely behind someone else who has used his or her credentials. Social engineers sometimes impersonate others, either in person or over the phone. Dumpster diving is the practice of sorting through trash to get potentially useful information that may have been thrown away. Pharming attacks redirect users to bogus websites.

A primary method of combating social engineering attacks is through training. The goal is to raise user awareness and change user behaviors from unsafe actions to safe actions. Training can be completed when a person is initially hired, annually, and/or periodically throughout the year. Training can consist of live presentations, online presentations, periodic e-mails, articles in newsletters, or any other means that an organization finds effective.

Questions

1. What is an APT?
 - A. A group, often sponsored by a government, that has the capability and intent to launch persistent attacks against an organization
 - B. Software that alerts a user that their system is infected with malware, but won't remove the malware unless the user pays a fee
 - C. An attack that redirects users to a bogus website
 - D. A scan to detect open ports
2. What is the difference between a DoS attack and a DDoS attack?
 - A. There is no real difference.
 - B. A DoS attack uses technical methods, but a DDoS attack uses nontechnical methods.
 - C. A DDoS attack is an attack from a single system, but a DoS attack is an attack from multiple systems.
 - D. A DoS attack is an attack from a single system, but a DDoS attack is an attack from multiple systems.
3. Of the following choices, what is a common DoS attack?
 - A. TCP flood
 - B. Tailgating
 - C. Smishing
 - D. Whaling
4. Thousands of computers have been infected with malware and are periodically directed to send out spam to other computers. What does this describe?
 - A. Zombies
 - B. Spear phishing
 - C. A botnet
 - D. Phishing
5. Which of the following best identifies a computer controlled by a botnet?
 - A. DoS computer
 - B. DDoS computer
 - C. Attacker
 - D. Zombie

6. An attacker is using Wireshark to capture and analyze TCP sessions. What is the best term that identifies this action?
 - A. Dumpster diving
 - B. Shoulder surfing
 - C. Sniffing
 - D. Vishing
7. A system has a protocol analyzer installed. What mode must the system operate in to capture all packets that reach it, including those that are not directly addressed to or from the system?
 - A. Promiscuous
 - B. Nonpromiscuous
 - C. DoS
 - D. DDoS
8. An attacker has written a program to shave off a penny from each transaction and divert the penny to the attacker's bank account. What best describes this attack?
 - A. Salami attack
 - B. Sniffing attack
 - C. Replay attack
 - D. Covert channel
9. Of the following choices, what provides the best protection against buffer overflow attacks?
 - A. SQL injection
 - B. Input validation
 - C. Cross-site scripting
 - D. Code signing
10. An application has received more input than it expected and the resulting error has exposed normally protected memory. What is the best explanation for what happened?
 - A. Phishing attack
 - B. Salami attack
 - C. Buffer overflow
 - D. Session hijacking
11. What type of attack can access data in a database used by a website?
 - A. Cross-site scripting
 - B. Cross-site request forgery
 - C. Rootkit
 - D. SQL injection

12. A website is preventing users from entering the < and > characters when they enter data. What is the website trying to prevent?
 - A. SQL injection attack
 - B. Cross-site scripting attack
 - C. Input validation attack
 - D. Trojan horse
13. A user receives an e-mail indicating that the bank has detected suspicious activity on the user's bank account. The message indicates the user should log on immediately to prevent loss of funds. What is the best term to describe this attack?
 - A. Sniffing
 - B. Session hijacking
 - C. Phishing
 - D. Tailgating
14. An attacker sends an e-mail to many members of an organization and spoofs the From address so that the e-mail looks like it came from within the organization. The e-mail tries to trick recipients into following a link. What is the best definition of this action?
 - A. Phishing
 - B. Spear phishing
 - C. Whaling
 - D. Vishing
15. A system has been attacked by an exploit that isn't published. What type of attack is this?
 - A. Scareware
 - B. APT
 - C. Pharming
 - D. Zero day
16. A user connected to a free wireless network at a coffee shop to access Facebook. Later, someone else started making posts on the user's page. What is the most likely cause of this?
 - A. Zero day exploit
 - B. WPS cracking
 - C. Evil twin
 - D. WPA cracking

17. An attacker uses nontechnical means to learn the e-mail address of a manager within a company. Which of the following best describes this attack?
- A. Social engineering
 - B. Shoulder surfing
 - C. Smishing
 - D. Covert cramming
18. Of the following choices, what is the best method to prevent tailgating?
- A. Education
 - B. Mantrap
 - C. Antivirus software
 - D. Access controls on the phone system
19. A user attempted to access <http://mcgraw-hill.com/> but was redirected to a website that advertised pharmaceutical drugs for sale. What does this describe?
- A. Phishing
 - B. Impersonation
 - C. Whaling
 - D. Pharming
20. What is a primary goal of security-related user awareness training?
- A. Increase use of e-mail
 - B. Change behavior
 - C. Implement technical solutions
 - D. Show how to use applications

Answers

1. **A.** An advanced persistent threat (APT) is a group of people (often sponsored by a government) that has the capability and intent to launch persistent attacks against organizations. Scareware is software that alerts a user their system is infected with malware, but won't remove the malware unless the user pays. Pharming is an attack that redirects users to a bogus website. A port scan is a scan that detects open ports.
2. **D.** A DoS attack is an attack from a single system, and a DDoS attack is an attack from multiple systems. Both typically use technical methods.
3. **A.** A TCP flood attack (also known as a SYN flood, TCP SYN, or TCP half-open attack) is a common DoS attack that withholds the third packet of the TCP three-way handshake. The other answers are not DoS attacks. Tailgating is a social engineering tactic. Smishing is a form of phishing using SMS messages. Whaling is a form of phishing against a single person, such as an executive.

4. **C.** A botnet is a group of computers that an attacker has taken over and now controls from a command and control center. The individual computers are referred to as zombies, but together they are a botnet. They may be directed to send out phishing or spear phishing e-mails, but that is the attack, not the network.
5. **D.** Computers controlled within a botnet are commonly called zombies. They are not referred to as DoS or DDoS computers, or attackers, although they can be directed to take part in a DDoS attack.
6. **C.** Sniffing is the practice of capturing and analyzing packets with a sniffer (a protocol analyzer). Dumpster diving refers to going through the trash looking for information. Shoulder surfing is the practice of looking over someone's shoulder to gain information, such as the password that a user enters to log on. Vishing is a form of phishing using telephones or VoIP.
7. **A.** The network interface card of the system running the protocol analyzer (or sniffer) must be in promiscuous mode. If it is in nonpromiscuous mode, the sniffer will only capture packets addressed directly to or from the sniffer. DoS and DDoS are not modes for a sniffer.
8. **A.** A salami attack uses multiple small, usually unnoticeable actions, such as shaving a penny off a transaction. A sniffing attack uses a sniffer (protocol analyzer) to capture and analyze traffic. A replay attack captures data and then later resends it to impersonate one of the parties. A covert channel uses an uncommon communications path to exchange information surreptitiously.
9. **B.** Input validation techniques validate data before using it and can help prevent a wide variety of attacks, including buffer overflow attacks. SQL injection is an attack that attempts to inject SQL code into an application. Cross-site scripting is an attack that attempts to inject HTML or JavaScript code into a web page. Code signing uses a certificate to digitally sign an application, but will not protect against buffer overflow attacks.
10. **C.** A buffer overflow occurs when an application receives more input than it expected and it is not able to handle the error gracefully. Attackers exploit buffer overflows to insert malware into systems. The best protection against a buffer overflow is to keep systems up to date. A phishing attack is sent through e-mail. A salami attack uses multiple small, usually unnoticeable actions, such as shaving a penny off a transaction. Session hijacking attempts to take over a session.
11. **D.** A successful SQL injection attack can access data in a database. Cross-site scripting injects HTML or JavaScript into a web page and runs the code on a user's system. A cross-site request forgery attack performs actions on behalf of a user without the user's knowledge. A rootkit is malware that takes over a user's system.
12. **B.** Cross-site scripting (XSS) injects HTML or JavaScript into a web page, and input validation techniques help prevent XSS attacks. The users are prevented from entering HTML or JavaScript tags that start with < and end with >. A SQL injection attack uses SQL code, but SQL code does not use < or > characters.

Input validation is a prevention technique, not an attack. A Trojan horse is an application that looks like it's something useful but is actually something malicious.

13. **C.** A phishing attack sends an e-mail to multiple recipients impersonating an e-mail from a legitimate company, indicating a problem, urging the recipient to take action, and warning of dire consequences if the recipient doesn't respond. A sniffing attack uses a protocol analyzer such as Wireshark to capture and analyze traffic. Session hijacking attempts to take over sessions and doesn't use e-mail. Tailgating is the practice of one person following another into a secure area while only the first person provides credentials.
14. **B.** Spear phishing is a phishing tactic that targets a specific organization. Phishing doesn't target individual organizations, but instead casts a wide net, hoping to catch someone. Whaling targets a specific individual, such as an executive. Vishing uses voice methods such as the telephone or VoIP.
15. **D.** Zero day exploits are attacks that take advantage of vulnerabilities that are unpublished and often include attacks that are unknown by the vendor. The other answers are known methods. Scareware is malware that scares users into thinking their system is infected with a virus and encourages them to install malware on their system. An advanced persistent threat (APT) is a group of people who have the capability and intent to launch extended attacks against organizations. Pharming is an attack that redirects users to bogus websites.
16. **C.** The most likely cause is an evil twin. An attacker likely created a free wireless hotspot in the coffee shop (perhaps on the attacker's laptop). When the user connected to it, the attacker captured the user's data, including logon credentials. This is a known attack, and whereas a zero day exploit is not widely known. Wi-Fi Protected Setup (WPS) cracking discovers the PIN of an access point and uses it to discover the access point's password. Wi-Fi Protected Access (WPA) cracking discovers the password on the access point by intercepting the four-way handshake and performing an offline brute-force attack.
17. **A.** Social engineering uses nontechnical (or low-technical) means to gain information, such as the names of people, e-mail addresses, and user credentials. Shoulder surfing is just looking over someone's shoulder, and although it may allow an attacker to see an e-mail address of a manager, it isn't the best answer. Smishing is a variant of phishing using SMS messages. There's no such thing as covert cramming.
18. **B.** A mantrap is the best method to prevent tailgating, which is the practice of one person following another into a secure area while only the first person provides credentials. Although education of employees can go a long way, ingrained courtesy sometimes overcomes security practices, and a person may actually open the door for a social engineer. Antivirus software, access controls, and the phone system aren't related to the social engineering practice of tailgating.

-
19. **D.** A pharming attack is one where the user is redirected to another website by manipulating one of the name resolution methods. Phishing involves sending an e-mail to many users and encouraging them to respond with personal information or by clicking a link. Impersonation, also known as masquerading or spoofing, is a social engineering tactic where the social engineer impersonates someone. Whaling is phishing attack that targets executives such as CEOs
 20. **B.** A primary goal of security awareness training is to change user behavior from unsafe practices to safe practices. It isn't related to applications such as e-mail, and end users aren't expected to implement technical solutions.

This page intentionally left blank

Malicious Code and Activity

In this chapter, you will learn about

- Different types of malicious code
- Common ways malware infects systems
- How attackers use remote access tools
- Stages of complex malware infections
- Differences between signature-based and heuristics-based antivirus software
- The importance of keeping antivirus signature definitions up to date
- The use of spam filters and content-filtering appliances
- How the principle of least privilege helps minimize infections
- Software security practices such as input validation and code signing
- Educating users about safe computing practices to reduce infections
- The Common Vulnerabilities and Exposures list

Identifying Malicious Code

Malicious code, or malicious software (malware), is any code or software that can be harmful or destructive to computers, networks, or the computing environment as a whole. Some of the common malicious code types include viruses, worms, logic bombs, Trojan horses, and ransomware.

Malware versions that are actively spreading and infecting computers are “in the wild.” In contrast, some researchers have discovered vulnerabilities in systems that attackers could exploit with malware. Researchers sometimes develop proof-of-concept code to demonstrate the potential exploits, but this code remains isolated within lab environments.

Virus

A virus is an application or a piece of code that causes unexpected and usually negative events on computers. Viruses have the ability to infect computers by copying themselves onto the computer using a wide variety of methods. Once the virus is on a computer, it replicates itself and at some point executes its payload, which is the unexpected or malicious event.



TIP Many people use the term *virus* when referring to any type of malware. In other words, someone may be talking about a Trojan horse, a worm, a logic bomb, or even spyware, but call it a virus. This is not technically accurate, because there are differences among these forms of malware, but the misuse of the term is still common.

One of the key characteristics of a virus is that the infected file must be executed for the virus to run. Typically, this means that a user must run a program to execute the virus. For example, imagine that Bob has an infected game file on his computer. If Bob starts the game, it executes the virus, but if Bob never starts the game, the virus never runs.

Users can execute a virus indirectly too. For example, many systems support an autoplay function. If a user inserts a CD, DVD, or USB drive into a system, the system automatically looks for an application and runs it. If this application is infected, the virus runs and can infect the system.



EXAM TIP A virus can only run with some type of user interaction. Typically, this entails a user running a program. However, in some cases, users can indirectly run a program by inserting a USB or optical drive into a system.

Viruses typically have two primary components:

- Replication component
- Activation component

A virus replicates itself by copying its code onto other applications. When a user runs the infected application, the virus runs and looks for other applications to infect. For example, a virus may wait for a user to insert a USB drive into a system. When the virus detects the USB drive, it copies itself onto a file on the USB drive. Similarly, when a user inserts the USB drive into another system, the virus tries to copy itself onto files on the new system.

At some point, the virus activates and delivers its payload, which is the malicious part of the program. The virus could delete or corrupt data or operating system files. If the virus is part of a botnet, it could check in with a botnet server for instructions. These instructions may be to download and install an additional virus, to send out spam, or even to launch an attack on another system. Chapter 5 discusses botnets in more depth.



NOTE Historically, viruses have been destructive to the installed system. However, with the explosion of botnets and criminalization of malware in recent years, this isn't often the case today. With the exception of ransomware, a virus will rarely destroy data on the system. Instead, it tries to stay hidden on the same system as long as possible. This allows the malware to continue to collect information from the user, or continue to use the computer as a zombie within a botnet.

Viruses come in many different forms based on what they attack or how they work. The following list isn't all-inclusive, but it does describe many of the methods and forms of viruses used over the years. It's worthwhile mentioning that many viruses combine different techniques to make them more difficult to detect.

- **Stealth** Many viruses try to hide themselves from antivirus (AV) software by providing false or misleading information about the virus to the AV software. For example, a file stealth virus can misreport its actual size so the AV scan thinks the infected file is the same size as the original. Additional methods used by stealth viruses include using armor, polymorphism, or metamorphism. You can compare these techniques to how a stealth aircraft hides itself from traditional radar and infrared detection devices, and remains hidden from enemies.
- **Armored** An armored virus uses code to make it difficult for AV researchers to reverse-engineer the code. When AV experts detect a virus, they attempt to decompile the code to determine what it does and how it operates. An armored virus uses different techniques to thwart the researchers' attempts. For example, an armored virus can use a compiler in combination with encryption to make it difficult to reverse-engineer the code. Developers often use complex and confusing code designed to hide what the virus is doing, along with encryption as a method of armor. Additionally, virus authors are constantly improving and inventing new armor techniques. An armored virus can also use stealth techniques to make it difficult for AV software to detect it.
- **Polymorphic** Some viruses have the ability to morph or mutate each time they replicate to another machine or even each time they run. Even though the file changes, the code used to replicate and deliver the payload remains the same. For example, by encrypting the file in a slightly different way each time the virus infects another file, it effectively changes the file. However, the virus replication and activation components will still behave in the same way. When the virus mutates, it's more difficult for AV software to detect it without additional virus signatures that recognize the mutated version. Virus authors use mutation engines to modify the viruses, and some viruses can mutate into thousands of different viruses.
- **Metamorphic** This is similar to a polymorphic virus, with a subtle difference. A metamorphic virus actually mutates the code used to replicate and deliver a payload. Although the function is the same, the steps used to do so are different, making the virus more difficult to detect. In contrast, a polymorphic virus will always perform the same steps to achieve the same function, while only the file changes.
- **Boot sector** The virus is stored on the hard drive's boot sector and executes when the system is booted. This includes both cold starts and reboots.
- **Multipartite** A multipartite virus has multiple components. For example, it could combine a boot sector virus with a virus that infects one or more files.

- **Macro** A macro virus resides in a document such as Microsoft Word, Microsoft Excel, or Adobe Reader. When a user opens the document, it runs the macro launching the virus. Most current versions of applications that support macros restrict their use to prevent the execution of malicious macros.



EXAM TIP You should be familiar with the different types of viruses, such as armored, polymorphic, and metamorphic, when preparing for the SSCP exam.

Worm

A worm travels over the network, looking for systems to infect. Worms are similar to viruses, but they do not require any type of interaction to execute. Instead, they slither and worm through the network on their own. If a worm infects a network, it has the potential to add a significant amount of traffic on the network and can effectively slow legitimate traffic to a crawl.



EXAM TIP A distinguishing difference between a worm and a virus is that a virus must be executed, while a worm can travel over a network automatically.

For example, a graduate student at Cornell University launched the Morris worm (the first known worm) in 1988 with the goal of gauging the size of the Internet. It searched for computers and installed itself on each computer it found and then looked for other computers. It often installed itself multiple times on the same computer, consuming resources, and eventually making the computer unusable. It resulted in estimated damages between \$10 million and \$100 million and prompted the creation of the first Computer Emergency Response Team (CERT) at Carnegie Mellon University.

Worms use protocols to travel over a network. For example, a worm can use TCP/IP to discover computers on the network, access network shares, and then use file transfer protocols such as FTP and TFTP to copy files and infect the target computers. The ability of worms to travel over a network provides additional justification for a defense in depth strategy that includes the following components:

- Host-based firewalls on all systems
- Up-to-date AV software on all computers
- Reduced attack surface by removing unneeded protocols



EXAM TIP Networks include firewalls and other security devices at the boundary to protect the internal network from Internet-based threats. However, malware might still make its way onto the network. Host-based firewalls, up-to-date AV software, and hardening techniques such as removing unneeded protocols help protect systems from these threats.

Trojan Horse

A Trojan horse is an application that looks like it's one thing but is actually something different. It presents itself as a useful or legitimate program, but includes a malicious component.

Its name comes from the Greek story of the nine-year Trojan War. Greek soldiers burned their camp and sailed away after leaving a large wooden horse in front of the gates of Troy. The residents of Troy thought the war was over and thought the horse (a sacred animal to them) was a gift. They ultimately rolled the wooden horse into their otherwise secure city and celebrated their good fortune. However, soldiers hid inside the wooden horse and the rest of the Greek army was poised to return. The soldiers emerged in the middle of the night, killed the guards, and opened the gates to the city, letting the rest of their army in. The Greeks sacked the city, ending the Trojan War. It's a simple message I learned from a mentor who taught me how to play chess many years ago: "beware of gifts."



EXAM TIP A Trojan horse appears to be one thing but is something different. If a user installs new software and then notices suspicious activity, the user very likely installed a Trojan horse.

As an example of a Trojan horse, one attack sent tens of thousands of e-mails to users that looked like it came from CNN and listed several titles of recent news stories. Users who clicked any of the titles opened a website page that popped up a small window. The window indicated that the user's system didn't have the current version of Adobe's Flash Player, so the page couldn't display the news story. It then prompted the user to install the current version. However, if the user clicked Yes to install the new version of Flash, it actually installed malware instead.



TIP Adobe gives simple advice to protect yourself. Only update Adobe products by going directly to the Adobe website and downloading updates directly from there. The same advice is useful for any software or updates. Don't trust a pop-up window, but instead go directly to the source website to download updates for any applications.

Another type of Trojan is scareware, described in the following section.

Scareware

Scareware (sometimes called rogueware) is malware that describes itself as antivirus software. A user may visit a website and then see a pop-up similar to that shown in Figure 6-1, which indicates that a virus has infected the user's system. However, every user who visits the website will see the exact same message. The website hasn't truly detected a virus. Instead, the pop-up attempts to trick the user into installing malicious software on their system.

Figure 6-1

Scareware pop-up trying to trick a user into installing malware



NOTE Although I created the pop-up window shown in Figure 6-1, it is similar to some pop-ups that users see when visiting some malicious websites. There isn't a Trojan virus named fakemeout.takemymoney. However, scareware does try to fake you out with the intention of taking your money.

If the user clicks the Fix Now button, it prompts the user to download and install the bogus antivirus application. After installing and running it, the bogus software indicates it is checking the user's system, but it's all a show. Unlike actual antivirus software that truly scans a system to look for malware, scareware does not look for malware. However, it always reports that it has discovered multiple instances of malware on the user's system. It then informs the user that the free trial edition only scans their system. The user must pay to purchase the full version to remove the malware.

Although the extortionary nature of scareware is malicious enough by itself, most scareware applications do more. For example, the downloaded application could join the system to a botnet controlled by a criminal. It could download and install a remote access tool that allows a criminal to take control of the user's computer, or install a keylogger to capture all of the user's keystrokes and send the data off to an attacker to use later. The possibilities are limited only by the imagination and technical abilities of the scareware authors. History shows us that they have a lot of imagination when it comes to stealing money from victims, and their technical abilities continue to increase.

What should a user do after seeing a pop-up window like this? Close the pop-up window and leave the site. Users who don't have antivirus software installed should go directly to a reputable site to get it, and not try to install it from one of these warnings. Many brands of antivirus software are available, and some reputable companies offer free versions.

LizaMoon Installs Scareware with SQL Injection Attacks

In 2010 and 2011, criminals attacked thousands of websites with SQL injection attacks dubbed LizaMoon. (Chapter 5 describes SQL injection attacks in more depth.) Successful attacks installed code that modified hyperlinks in databases and redirected users to malicious websites.

When users clicked the malicious links, the target website popped up a warning from a fake “Windows Stability Center” in an attempt to impersonate a valid Microsoft product. The warning indicated user computers had security problems and urged users to purchase software to fix the problem.

Ransomware

Ransomware is a type of Trojan that takes control of the user’s computer and then demands a ransom from the user to get control back. There are many variants, but there are two primary themes. One type locks you out of your computer, preventing you from logging on. Another type encrypts your data files, allowing you to log on but not access any of your data.

The “Mandiant U.S.A. Cyber Security Ransomware” is an example of ransomware that locks you out of your computer. When users turn on an infected system, it displays a bogus but official-looking screen that indicates it’s a joint effort between Mandiant Cyber Security, the U.S. FBI, the U.S. Department of Defense, and the U.S. Cyber Crime Center. Variants (known as the Police Virus, Trojan Reveton, and Moneypak Virus) replace the names of U.S. agencies with those of agencies in other countries, such as the International Criminal Police Organization (Interpol). This type of malware isn’t a joint effort with any official agencies. Instead, it’s an attempt by criminals to extort money from users.



NOTE Information security company Mandiant has done a lot of research on attackers. It released a report in 2013 documenting activities of advanced persistent threats (APTs) operating in China. Mandiant was later purchased by FireEye, another security company.

The ransomware message accuses the user of participating in illegal activity such as distribution of pornography, copyrighted files, or malware. It then demands that the user pay a fine of \$300 (or some other amount) within 48 hours or they will face legal prosecution. The software directs users to purchase a MoneyPak or MoneyGram voucher at a store such as Kmart or Walgreens. They enter the voucher code into the ransomware screen to pay the attackers, which unlocks their computer.

CryptoLocker is an example of ransomware that encrypts user files. When it infects the user’s system, it locates and encrypts data files on the user’s system, attached USB drives, external hard drives, shared network drives, and sometimes cloud storage drives. It then demands payment of \$300 (or more) to get the decryption key. The message indicates that

if the user doesn't pay within 72 hours, the attackers will delete the decryption key, causing the user to lose the files forever. Attackers sometimes offer to retrieve the decryption key after the time limit has expired, but at a much steeper price as high as \$2,300. Victims paid approximately \$30 million between September and December 2013.

Police forces disrupted the network running the original CryptoLocker ransomware in June 2014. However, CryptoLocker and many variants began appearing soon afterward. TorrentLocker is similar and often uses the CryptoLocker name, but the underlying code is different. The FBI is offering a \$3 million reward for information leading to the arrest or capture of a Russian man (Evgeniy Mikhailovich Bogachev) who they believe is the ringleader of the gang that ran the original CryptoLocker network.



TIP FireEye partnered with Fox-IT to help users decrypt files encrypted by CryptoLocker. Users send an encrypted file with their e-mail address. The FireEye/Fox-IT software attempts to discover the master decryption key and e-mails it back to the user. Users can then decrypt files on their system using the master decryption key and a free download available through the same site: <https://www.decryptcryptolocker.com/>. Unfortunately, this might not work with some new versions of CryptoLocker.

Notice that both of these ransomware methods include a sense of urgency. The Mandiant virus states that users will be prosecuted if they don't pay the fine within 48 hours. CryptoLocker states that attackers will delete the decryption key if users don't pay the ransom within 72 hours. Both types include a timer that counts down as a constant reminder to the user to take action.

Installing CryptoLocker

Criminals have used a variety of methods to trick users into installing CryptoLocker. Many methods include tricking users via fake e-mails. Many fake e-mails have mimicked UPS and FedEx tracking notices indicating the delivery service attempted to deliver a product but failed. These notices encourage the user to either click a link to a (malicious) website or open an (infected) attachment. Visiting the website or opening the attachment launches the attempt to install the malware.

As another example, attackers have sent to users e-mails that looked like they were from an official state agency. The e-mails said that a traffic camera caught the recipients speeding and that they now had to pay a fine to avoid the suspension of their driver's license and/or vehicle registration. The e-mails included two buttons, one for an invoice and one to view the camera images. Clicking either button took the recipients to a web page that prompted them to enter a CAPTCHA code. When they did, it downloaded a file. If they opened the file, it installed CryptoLocker and encrypted all the data on their system.

Keylogger

A keylogger (also known as a keystroke logger) is a tool that can capture all the keystrokes that users enter on a computer. Keyloggers can be hardware based or software based. A hardware-based keylogger plugs into the computer between the keyboard and the computer and captures the data. For example, if the keyboard plugs into a USB slot, you would plug a USB-keylogger into the USB slot, and then plug the keyboard into the USB-keylogger. It could transmit the information wirelessly or store the data on the hardware.

Software-based keyloggers are more common. They simply run as an additional piece of software when the computer starts. For example, some parents use keyloggers to monitor children's computer activity. Keyloggers can monitor online chats, instant message content, e-mails, and data entered into Internet website forms. Many keyloggers also include the ability to capture screenshots. These same tools are available for employers to monitor employee activities.

Of course, criminals use the same type of keylogger software to capture information entered by users. Users don't knowingly install the keylogger. Instead, another infection such as a Trojan horse installs the keylogger.

Logic Bomb

A logic bomb is malicious code that executes in response to an event such as a point in time or in response to specific actions. For example, Chapter 5 mentioned the attempted attack on Fannie Mae by a terminated administrator. After his employer terminated him from the job, he added a script that was set to run on January 31 that would have deleted passwords, data, and backups for about 4,000 servers. This is a logic bomb because it was set to run on a specific date.

Some keyloggers have a logic bomb component. They send data to an attacker after a user accesses specific websites, such as banking sites. For example, when the keylogger detects that the user accessed mybank.com and entered a username and password, the keylogger e-mails the URL and credentials to an attacker.

Rootkits

A rootkit is a set of programs that can run on a system largely undetected. It has root-level access to the system, similar to how a root-level administrator has full and complete control over a system. Attackers can often control the computer remotely after installing a rootkit, and the user typically doesn't notice any symptoms.

Rootkits have low-level control on systems, which allows them to identify when anti-virus software is running. The rootkit can detect a system scan and hide itself from the scan or restrict the data returned to the scan. This has made rootkit detection difficult for traditional AV software. However, it's worth noting that attackers typically install rootkits via a virus or a Trojan. In other words, if you can detect and prevent attackers from installing viruses and Trojans on your system, you have a good chance at preventing them from installing rootkits.

Several companies have released rootkit detectors in the last few years. For example, Sophos has released a free dedicated tool designed to detect and remove rootkits from systems. If interested, you can read about it here: www.sophos.com/en-us/products/free-tools/sophos-anti-rootkit.aspx. Additionally, most popular AV programs advertise their ability to detect and remove rootkits. This typically requires rebooting the system from a CD or bootable USB with antivirus software to detect the rootkit and clean the system.



TIP Microsoft 64-bit systems have added more protections against rootkits. They use the Kernel Patch Protection (also known as Kernel PatchGuard), which prevents rootkits from modifying the underlying operating system components (the kernel). Kernel PatchGuard requires digitally signed drivers. This ensures the integrity of the files. If malware modifies the files, the digital signature will detect the modification.

Mobile Code

Mobile code is any code that can transfer itself between systems without the user taking explicit action to install it. Attackers use malicious mobile code to install malware from remote servers via the Internet. Malicious mobile code is typically transmitted via interactive web applications. Some examples of mobile code include the following:

- **Scripts such as JavaScript or VBScript** Attackers embed the script into a web page, or within any file, that supports the script. When users click a link to retrieve a web page, the web server sends the page along with the embedded script. As the web browser displays the web page, it also runs the script. Some HyperText Markup Language (HTML)-based e-mail applications support the use of scripts, so opening an e-mail can launch the script. Similarly, files such as Portable Document Format (PDF) documents and Microsoft Office documents support scripts. Attackers embed the code into these documents, which executes when the user opens the document. While keeping systems and applications patched with security updates protects against most of these threats, many users don't keep their systems patched.
- **Java applets or ActiveX controls** Java applets and ActiveX controls are mini programs that developers embed into web pages. Similar to scripts, they run when the web browser displays the web page. A primary method of protection is to block unsigned applets and controls. Using code signing (as discussed in Chapter 5) is one method that developers use to show their code is valid and provides assurances to users that the code is valid.
- **Documents that include macros** Many document types (such as Microsoft Office documents) support macros. While macros can be very helpful, an attacker can write malicious mobile code into a macro. When a user opens the file, it runs the macro. To reduce this risk, it's common to disable macros in applications that support them.

Backdoors and Trapdoors

A trapdoor or backdoor is code that is embedded in an application and provides access to the application, the application's code, or data via a covert method. Some backdoors provide access to the operating system.

Application developers often add backdoor code to an application for debugging purposes. For example, by entering specific keystrokes, they can display the value of variables within a running application. Although this is useful during the debugging process, a backdoor can be dangerous if developers leave it in the released version of the application. Instead, as a best practice, most developers remove all backdoor code prior to releasing the application.

Attackers often attempt to install a backdoor after infecting a system with a virus or Trojan. When successful, attackers can use backdoors to read data, install additional malware, or even take control of the user's computer from a remote location. Within the context of malicious code, both backdoors and trapdoors mean the same thing, although the term *backdoor* is more common.

Some attackers have successfully added backdoor code to released applications. For example, the Mydoom worm installed a backdoor that affected versions of Microsoft Windows and allowed attackers to control the infected computers remotely over the Internet. As with most other malware, after security experts discovered the vulnerability, the vendor (Microsoft in this example) wrote and released security updates to protect systems from the threat.



NOTE The Mydoom worm (released in 2004) was the fastest-spreading e-mail worm up to that date. It infected computers and then used them to send spam. The worm created a backdoor on port 3127, which allowed remote control of the infected PC. Variants appeared in 2005 and then again in 2009.

Another example is the Win32/Zbot (also called GameOver Zeus or simply Zeus) family of password-stealing Trojan horses. This is also a good example of a multipartite virus with backdoor functionality. Win32/Zbot tricked users into installing the Trojan onto their systems, and then joined their computers to a botnet. Interestingly, infected systems don't join a single huge botnet, but instead join one of many independently operated botnets. The attacker controlling the command and control center for each of these botnets then controls these infected systems through the botnet's backdoor.

Win32/Zbot successfully used a variety of techniques to steal financial-related information such as usernames and passwords. For example, it interfaced directly with web browsers to monitor traffic and information. The Trojan deleted cookies for websites, forcing users to reenter their credentials. It injected additional HTML code into banking website pages, requesting additional information from users such as Social Security numbers or birthdays. Win32/Zbot then encrypted this information and sent it to attackers controlling the botnet. Criminals logged on to the victim accounts with the stolen credentials and transferred funds to other bank accounts controlled by the criminals.



NOTE The FBI believes that the ringleader of the network that released the original CryptoLocker ransomware also created and distributes the Zeus banking Trojan. Security professionals estimate that attackers have successfully used Zeus to steal hundreds of millions of dollars. There are also indications that Zeus malware installed CryptoLocker onto infected systems.

RATs

A remote access Trojan or remote access tool (RAT) provides an attacker with control over a target computer via the Internet. Many Trojans install RATs, and then attackers can use the RAT to do anything on the target computer. This includes using the RAT to install malware, install keyloggers, and join the compromised computer to a botnet.

Spyware

Spyware is software that can install itself on a user's system without the user's knowledge or consent. Although there was a time when spyware was relatively benign, it is much more dangerous today. For example, there was a time when spyware attempted simply to understand a user's surfing or buying habits so that it could send targeted advertisements to the user. However, spyware today is much more malicious and attempts to learn details to steal financial data, logon information, or a person's identity.

Most security experts consider spyware a form of malware because it can cause so much damage to an individual. Many spyware applications can identify websites that a user visits, usernames and passwords that a user enters, user's keystrokes even when not visiting a website, and even search a user's computer for information such as files including usernames and passwords. Spyware can also change settings on a user's system, such as settings to enable or disable a firewall, settings to enable or disable automatic updates for the operating system, or even settings for individual applications.

Due to the malicious nature of spyware today, most AV software includes the ability to search for and detect spyware. Years ago, many users used one application for spyware and another application for viruses and other malware. However, if a user gets an AV application that includes a spyware component, it isn't necessary to install a separate spyware application.



NOTE An advanced type of spyware is canvas fingerprinting, which uses a variety of techniques to track users through an entire surfing session. A large-scale study titled "The Web never forgets: Persistent tracking mechanisms in the wild" documents three advanced web tracking mechanisms. You can read about it here: <https://securehomes.esat.kuleuven.be/~gacar/persistent/index.html>.

Malware Hoaxes

In its simplest terms, a malware hoax is a false message about a malware risk that doesn't exist. Attackers and uneducated users typically spread hoaxes through e-mail while urging other users to forward it to everyone they know. Many times these e-mails encourage users to take action that can harm their system.

Hoaxes commonly include a sense of urgency and often overuse capital letters and exclamation marks, with subject lines featuring something like “URGENT!!!,” “WARNING!!!,” or “VIRUS ALERT!!!” In contrast, legitimate warnings (such as those from a cyber security alert service or from an organization’s administrator) simply provide the detailed information without stimulating emotion or fear.

In general, hoaxes just waste people’s time. When users forward it to everyone they know, they perpetuate the hoax, along with consuming more bandwidth and wasting other people’s time. Many organizations provide training to users so that they don’t forward e-mail warnings to everyone. Instead, users are encouraged to forward these types of e-mails to system administrators if they think it represents a real threat. Although this adds to the administrator’s workload, it’s better than spamming the rest of the organization.

Analyzing the Stages of Regin

While the previous descriptions of malicious code describe individual components of malware, they really don’t paint the full picture of malware threats. Symantec released a white paper on Regin (an advanced class of malware) in November 2014, which demonstrates the complexity behind some of the serious malware threats. The ultimate goal of Regin is to collect data, and it does so through several different methods, including using Trojan RAT features, capturing screenshots of user computers, taking control of the mouse, stealing passwords, monitoring network traffic, retrieving deleted files, and collecting e-mail.



TIP This section describes the known stages of Regin. However, it’s important to remember that this is just one example. Many families of similar malware include the same goals with the same level of complexity.

Regin 1.0 has been around since at least 2008, and Regin 2.0 has been used since at least 2013. This speaks to the technical sophistication of the attackers. Developers aren’t teenage script kiddies doing random attacks. Rather, the developers appear to be a cohesive team with specific goals of monitoring targets and stealing data. There are six distinct stages to Regin, as shown in Figure 6-2 and described in the following list:

- **Stage 0** A dropper is a type of Trojan horse that attempts to drop malware onto a computer. Within Regin, stage 0 is the dropper stage and represents the initial infection. It could be from the user opening an infected file in an e-mail, clicking a malicious link in an e-mail, or any other way users inadvertently install Trojans. Stage 0 installs and executes stage 1 files. In many cases, stage 0 code is only stored in memory and disappears after stage 1 starts.
- **Stage 1** This stage loads operating system kernel drivers when the computer starts. These drivers decrypt, load, and execute stage 2. Two known names for these files are `subclass.sys` and `adpu160.sys` for version 1 and version 2, respectively. Stage 1 is the only stage that includes plainly visible code on the attacked computer. Regin encrypts stages 2, 3, 4, and 5 and stores them in files,

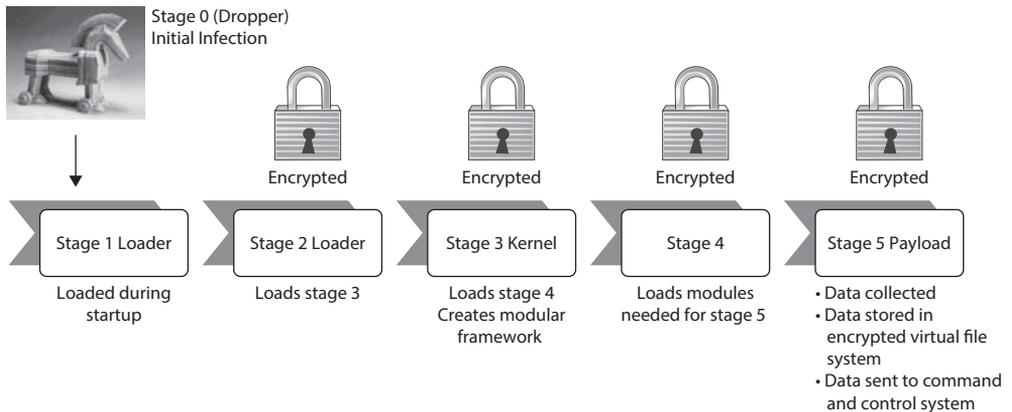


Figure 6-2 Regin stages

or within nontraditional areas such as the registry, or within raw sectors of the disk drive. As mentioned previously, encryption is a method of armoring the virus to prevent AV researchers from decompiling and analyzing the code.

- **Stage 2** This stage loads another operating system kernel that decrypts, installs, and runs stage 3. Regin stores encrypted data for this stage within an NTFS extended attribute, or in some cases within one or more registry subkeys. This stage often hides running instances of stage 1, removing any plainly visible elements of Regin.



NOTE NTFS extended attributes are name–value pairs used for a variety of purposes. For example, Windows systems use them to mark files downloaded from the Internet, and provide a warning to users when they open these downloaded files. Regin uses the %Windir%, %Windir%\fonts, and %Windir%\cursors extended attributes to store encrypted blobs. You might know these as environment variables, which they are. However, Windows systems store environment variables as extended attributes.

- **Stage 3** This stage decrypts, installs, and runs stage 4 files and provides a modular framework for stages 3, 4, 5, and 6, consisting of hundreds of programmatic methods. These modules provide a wide range of functions such as compression and decompression routines, encryption and decryption routines, an encrypted virtual file system, and networking routines to send data to the attacker. Stage 3 is about six times larger than stage 2.
- **Stage 4** Stage 4 loads the relevant modules needed for stage 5 depending on the targeted payloads. Because Regin includes so many different modules, attackers

can modify it to collect custom payloads. For example, one implementation might collect e-mail, another implementation might install a RAT, and so on. Note that this isn't a different version of Regin, but instead, the attacker just enables different features of Regin by implementing different modules.

- **Stage 5** Stage 5 collects data from a variety of sources and stores that data within the encrypted virtual file system. It can collect network data, passwords, computer files, deleted files, and screenshots. Stage 5 sends data to the attackers via one of several standard protocols, using special flags for identification.

While these six stages occur automatically after the initial infection, Regin also supports command and control operations similar to a traditional botnet. Attackers can contact the computer periodically and instruct it to send data. This often occurs using two stages. For example, the attacker might send an Internet Control Message Protocol (ICMP) datagram with instructions to send collected data to a specific server. The infected computer then responds with the payload using User Datagram Protocol (UDP) datagrams. Regin encrypts all data before sending it to the command and control computer.

Understanding Malware Delivery Methods

Attackers deliver malware through multiple methods. While most malware is delivered via the Internet, it can also be delivered from hardware devices such as USB drives. When end users understand how malware spreads, they are less likely to be tricked into installing malware on their systems.

Delivering Malware via Drive-by Downloads

In a *drive-by download*, web servers include malicious code that attempts to download and install itself on user computers when the users visit. These drive-by downloads can include many different types of malware, such as installing a RAT or joining a computer to a botnet.

Criminals rarely host these computers themselves because that would make it easier for law enforcement to find them. Instead, they look for vulnerable computers that they can modify to include the drive-by downloads. For example, buffer overflow attacks (described in Chapter 5) give attackers access to a web server's memory and allow the attackers to install malicious code.



TIP Attackers modify vulnerable servers to install other malware too. For example, scareware uses a pop-up to indicate the site detected malware on the user's computer and then attempts to trick users into installing fake antivirus software. It's easy to think that the criminals own the site, but more often attackers just install their own pop-up on an innocent business owner's site.

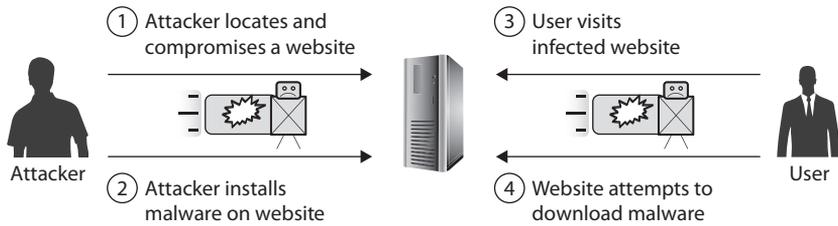


Figure 6-3 Drive-by download installing malware

Here are the typical steps involved in a drive-by download, depicted in Figure 6-3:

1. Attackers compromise a website to gain control of it.
2. Attackers install a Trojan or other malware embedded in the website's code.
3. Attackers attempt to trick users into visiting the site. Sometimes, they simply send the link to thousands of users via e-mail, hoping that some of them click the link. Other times the link is from another website or a malicious banner ad.
4. When users visit, the website attempts to download the Trojan or malware onto the user's system.

Users can prevent successful drive-by downloads by keeping their systems up to date with operating system and application updates and with up-to-date antivirus software. This will prevent most infections, but it won't prevent zero day exploits, so users still need to be vigilant. A *zero day exploit* is an attack that takes advantage of unpublished vulnerabilities that vendors either haven't heard about or haven't published updates for yet.

Delivering Malware via Malvertising

Attackers sometimes deliver malware via malicious banner ads, commonly called malvertising. These look like regular ads, but they contain malicious code. Many of them include a link taking users to a server hosting a drive-by download. Others are Flash applets with malicious code embedded in them.

Malvertising has appeared on mainstream websites such as those belonging to The New York Times and Yahoo!. As an example, ads hosted on Yahoo! in 2013 and 2014 took users to sites hosting scareware. They included fake pop-ups as described in the "Scareware" section earlier in this chapter. In mid-2014, ads on Yahoo! took users to Eastern European sites hosting CryptoWall ransomware, a clone of CryptoLocker.

It's worth noting that legitimate sites are not knowingly hosting these malvertisements. Instead, they sell ad space to advertising companies and these advertising companies unknowingly send the malicious ads to the legitimate sites. In some cases, criminals trick legitimate advertising companies into hosting the ads. For example, social engineers tricked people at the Gawker Media network to sell fake Suzuki advertisements. In other cases, attackers compromise websites and post their own malicious ads on the websites.

Delivering Malware via E-mail

One of the most popular methods of delivering malware is via e-mail. Chapter 5 discusses spam, botnets, and various types of phishing attacks. As a reminder, criminals control hundreds of thousands of compromised systems in botnets and they rent these botnets out to other criminals who use them to send phishing e-mails.

While some phishing e-mails attempt to trick users into giving up personal information such as their credentials for a banking site, other e-mails try to install malware on the user system.

One method includes infected attachments. If users open the attachment, it can infect the user's system. Attachments arrive as ZIP files, PDF documents, web pages (as HTML files), apparent graphics such as JPG files, and more. One way they disguise these documents is to simulate the extension of a safe file. For example, a JPG graphic file has a .jpg extension, such as picture.jpg. An attacker can rename an executable file to picture.jpg.exe. Because the operating system normally hides the extension, the file looks like picture.jpg, indicating it's a relatively safe JPG file.

It's worth noting that attackers often spoof or modify the From e-mail address to make it look like it's from someone the recipient knows. For example, attackers monitor social media sites and identify your friends. They then send an e-mail to you but change the From field so that it looks like it came from your friend. This tricks users into trusting the e-mail and makes them more willing to download an attachment or click a link.

Delivering Malware via USB Drives

As mentioned earlier in this chapter, some viruses replicate via USB drives. An infected USB drive will try to infect a computer system when a user inserts it. An infected system waits for a user to insert a USB drive and attempts to infect the drive.

When an organization allows users to transport data back and forth from home and work using USB drives, it becomes very easy for a user to transport viruses from home and work. Although an organization can implement multiple methods to protect its environment from malware, the organization usually has no control over a user's home network. If users are authorized to use USB drives, it's imperative that each system have AV software installed and kept up to date. Some organizations use additional methods to prevent systems from reading data from a USB drive, writing to a USB drive, or both.



EXAM TIP If users are authorized to transport data with USB drives, each system should be protected with AV software. This is the best method to protect a system without affecting the user. It's also possible to prevent copying of data with USB drives to provide stronger protection.

Implementing Malicious Code Countermeasures

Because malware is so prevalent, it's important to protect systems from malware. The primary method of protection is antivirus software and an aggressive defense in depth

strategy. As an overview, the following methods are effective countermeasures against malware:

- Install antivirus software on all systems.
- Install antivirus and antispam software on e-mail servers.
- Install antivirus and content-filtering software on firewalls.
- Keep all antivirus and antispam software up to date.
- Perform regular antivirus and vulnerability scans.
- Keep all systems up to date.
- Educate users.

Antivirus Software

The primary method used to detect and prevent infections is the use of AV software. In addition to being able to detect and prevent infections, AV software can remove the malware, restore infected files, or quarantine infected files if the AV software can't remove the malware. The goal is to ensure that the malware cannot run and cause any more problems on a system. Researchers can examine and analyze quarantined files. However, the computer will not run quarantined files.

Figure 6-4 shows an example of antivirus software. It includes the ability to perform on-demand scanning (shown as Scan Your PC in Figure 6-4), real-time scanning, and scheduled scanning:

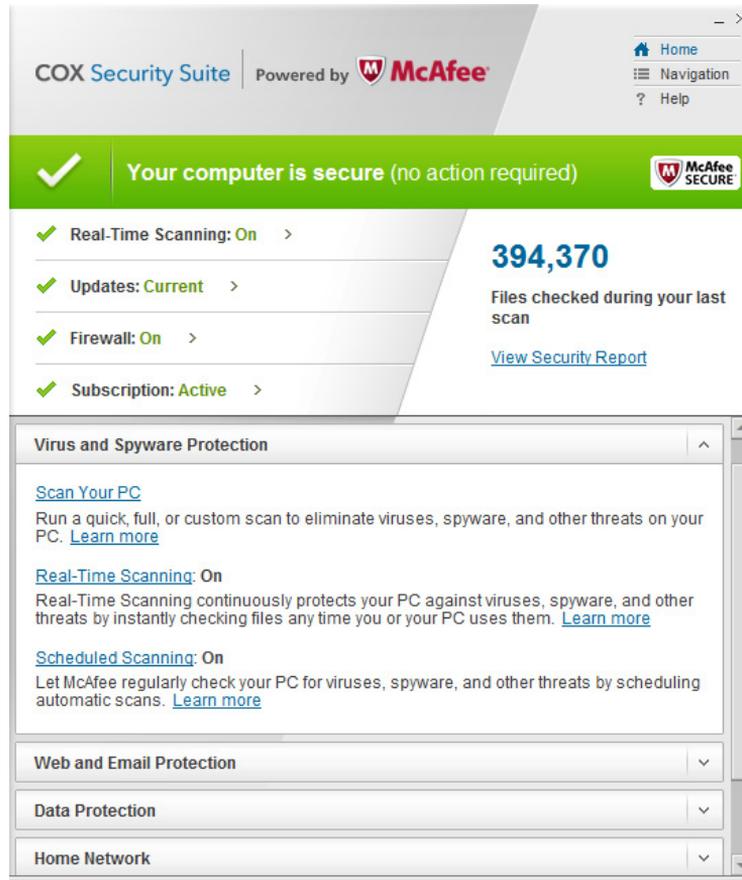


NOTE Many Internet service providers (ISPs) provide free copies of antivirus software for their customers. For example, at the top of Figure 6-4, you can see the COX Security Suite branding. Cox Communications provides this software to its Internet customers.

- **Real-time scanning** This provides protection as the user opens and uses files. For example, if a user opens an e-mail attachment, the AV software checks it to ensure that it doesn't have any malicious content. Real-time scanning checks any file that is opened or launched by either a user or the operating system.
- **On-demand scanning** If you suspect a computer is infected, you can perform an on-demand scan to check it. Most AV software allows you to run a scan to check memory, specific files, folders, or drives, or complete a full scan of the entire system. Figure 6-5 shows an on-demand scan in progress on a system.
- **Scheduled scanning** A scheduled scan occurs on a regular basis, such as once a week. Once scheduled, the AV software launches the scan based on the schedule.

Antivirus software uses a software engine to detect viruses using virus signatures. Additionally, some antivirus software uses heuristic scanning. The following sections describe both methods.

Figure 6-4
Antivirus
software



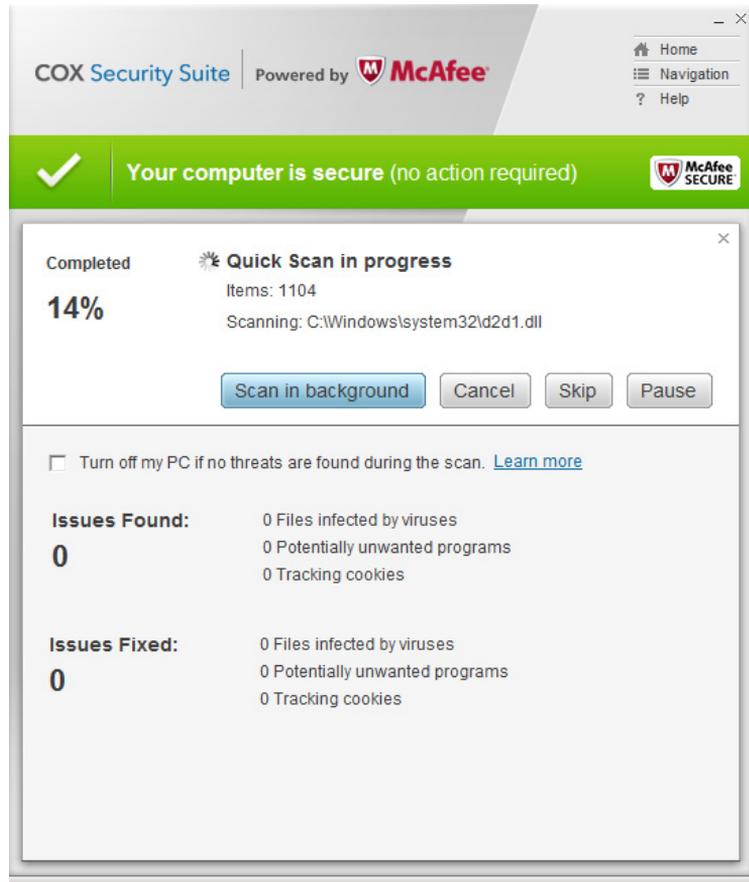
Signature-based Detection

Viruses have specific characteristics that AV software uses to identify them. The signature can be a unique characteristic such as a specific byte pattern within the virus. This is similar to how a person's fingerprint can uniquely identify an individual. Viruses in files, memory, and a drive's boot sector all have specific signatures that AV software uses to detect them.

When AV researchers discover malicious software, they identify its signature. AV companies add this signature definition to the AV signature database, which the AV software uses to detect the malware. When the AV software finds a file that matches a signature in the database, it alerts the user and often prompts the user to remove, repair, or quarantine the file.

You may remember from earlier in the chapter that polymorphic and metamorphic viruses change. A primary goal of attackers using these techniques is to prevent detection of the virus through signatures. Additionally, attackers have developed application tools

Figure 6-5
Performing an
on-demand scan



that allow them to create new iterations of viruses rather easily. For example, the attacker can write the replication and activation component to create the core virus. The attacker can then run the virus through another application to create hundreds of different versions of the same virus. Some viruses can generate as many as 500 different variants in a single month. In other words, if a virus has a specific signature but it has morphed into a different variant, it's possible that the original signature will no longer discover the virus.

Of course, as viruses morph to a state that prevents detection, AV researchers create signatures to detect the new variants. As long as users ensure their AV signature files are updated regularly, the AV software will discover these new variants. However, if users don't update the virus signatures regularly, the new variants may run undetected.

Heuristic-based Detection

Attackers release new viruses into the Internet regularly, and older viruses are regularly mutating into variants. Because of this, it's entirely possible that at any time, a system will come across a previously unknown virus. This means that the AV software doesn't

Staying Aware of Malware Threats

Malware is very dynamic. As an example, several international law enforcement agencies disrupted the crime ring that was running the GameOver Zeus botnet and the CryptoLocker ransomware in mid-2014. Many experts said this was the end of this criminal ring and the end of CryptoLocker. However, CryptoLocker and some variants reappeared a few months later, and are now as strong as ever. Additionally, the leader of that gang remains so active that the FBI has offered a \$3 million reward for information leading to his arrest and conviction. One way to keep up with these changes is to subscribe to various newsletters or bulletins. Here are a few that you might find useful:

- **Virus Bulletin** This organization focuses on preventing, detecting, and removing computer threats, including malware and spam. They previously published a monthly bulletin, but now publish them at least monthly, and more often when appropriate. You can view recent or archived articles on their website (<https://www.virusbtn.com/>) and it's free.
- **SANS Internet Storm Center** This site is part of the SANS Technology Institute and it provides information on current threats. You can sign up to receive daily podcasts at <https://isc.sans.edu/dailypodcast.xml> and view past diaries at <https://isc.sans.edu/diaryarchive.html>.
- **Cyber Awareness Alerts** These alerts provide timely information about current security issues, vulnerabilities, and exploits. They come out at least once a month and provide technical information for system administrators and experienced users. You can view archives here: <https://www.us-cert.gov/ncas/alerts>.
- **Cyber Awareness Bulletins** These provide weekly summaries of new vulnerabilities. They come out once a week and include information about security issues and vulnerabilities from the previous week. You can view archives here: <https://www.us-cert.gov/ncas/bulletins>.
- **Cyber Awareness Tips** These tips provide information and advice on common security topics. They are written for regular users and come out every two weeks. You can view archives here: <https://www.us-cert.gov/ncas/tips/>.
- **Current Activity Updates** These provide timely updates about high-impact security activity. This list can generate alerts from several times a day to several times a week. You can view archives here: <https://www.us-cert.gov/ncas/current-activity/>.

You can sign up for any (or all) of the Cyber Awareness notifications here: <https://public.govdelivery.com/accounts/USDHSUSCERT/subscriber/new>.

have a signature for the new virus. However, heuristic scanning can often detect these previously unknown viruses.

Heuristic-based detection attempts to detect a virus based on its behavior. In other words, if a file is behaving in a way that is not typical for a file, the heuristic engine can detect this behavior, identify the file as malware, and block the malware from infecting the system. For example, legitimate applications rarely modify, copy, or delete system files. If an application does attempt to take one of these actions, the heuristic-based engine can detect and block it.

Many heuristic scanners use a sandbox technique, which is an isolated or confined area of memory on the system. The AV software ensures that new or modified applications run within an isolated sandbox for analysis before allowing them to run regularly. The idea is that the application can play within its own sandbox, but it can't interfere with applications or processes in other sandboxes or other areas of memory. The scanner observes the behavior of the application as it runs within the sandbox. As long as it doesn't display any suspicious activity, the AV software allows it to run.



EXAM TIP Virus signatures detect known viruses. It's important to update AV software signatures regularly to ensure the AV software can detect newly released malware. Heuristic-based detection attempts to detect previously unknown viruses.

Keeping AV Signatures Up to Date

An important element of using AV software is to keep the signature definitions up to date. Because viruses can morph so quickly, most AV software includes the ability to check automatically for new signatures, and they do so once a day and sometimes as often as once every hour. If the AV software is not up to date with current signature definitions, systems quickly become more vulnerable to new attacks.



TIP AV software downloads signature files incrementally so that they include only differences. In other words, if you have all the current signatures except for two, the download only includes the two new signatures, not the entire signature definition file.

Users often don't understand the risks and the importance of updating signature definitions, so many organizations typically take control of this step. For example, some AV vendors sell server products where a central server can deploy both AV software and updates to all the clients within an organization. Administrators ensure the central server downloads the updates from the vendor site and then deploys the updates to clients within the organization. Many times these updates include more than new signature definitions. They also upgrade and improve the antimalware application.

Spam Filters

Because attackers often deliver malware and malicious links via spam, a spam filter can provide an important layer of protection. A spam filter attempts to detect unsolicited e-mail and block it. Many e-mail servers include antis spam software to filter spam. Additionally, many e-mail applications include antis spam software to filter spam.



NOTE It's estimated that as much as 70 percent of all e-mail is spam. It clutters inboxes, consumes bandwidth, and delivers malware and malicious links.

A danger of a spam filter is blocking legitimate e-mail. For example, a company that sells products and communicates to customers via e-mail doesn't want to block legitimate queries from customers. Because of this, most spam filters have low thresholds for spam. Given the choice of accepting some spam but never blocking legitimate e-mail, or blocking more spam but occasionally blocking legitimate e-mail, most organizations choose the former. They'll allow some spam to get through but do not want to block any legitimate e-mail.

Sender Policy Framework (SPF) records are a newer technology used to reduce spam by identifying spoofed e-mail. SPF records identify e-mail servers authorized to send out e-mail for a domain. For example, McGraw-Hill can publish an SPF record to DNS servers on the Internet that identifies the mail.mcgraw-hill.com server as the only authorized server to send out e-mail for the mcgraw-hill.com domain. Spammers sometimes spoof the source address, but the e-mail headers show the server that actually sent the e-mail. Using SPF records helps detect some of this spam.

E-mail servers that receive e-mail can use SPF records to validate e-mail. When an e-mail server receives e-mail that appears to come from McGraw-Hill, it can check what server actually sent it and compare it to the public SPF record. If a different server sent the e-mail (not one identified in the SPF record), the receiving server marks the e-mail as spoofed. The server can then block all spoofed e-mail from this server to reduce spam levels. RFC 4408 formally defines SPF, although it is still in experimental stages.

Content-Filtering Appliances

A content-filtering appliance is a device that filters traffic in and out of a network. It actively monitors data streams by inspecting datagrams, searching for malicious code or malicious behavior. Figure 6-6 shows the overall concept. You install the device so that all traffic going to or from the Internet goes through the device. Basic content-filtering appliances filter malware and spam. More advanced appliances can act as firewalls, proxy servers, and intrusion detection systems. Notice that both malware and spam enter the appliance, but the device can block malware and spam, preventing it from getting into the internal network.

Similar to kitchen appliances, content-filtering appliances are rather easy to plug in and use. Think of a microwave kitchen appliance. It certainly has a lot of complexity within it, but all you have to do is plug it in, put your food in, punch some buttons, and

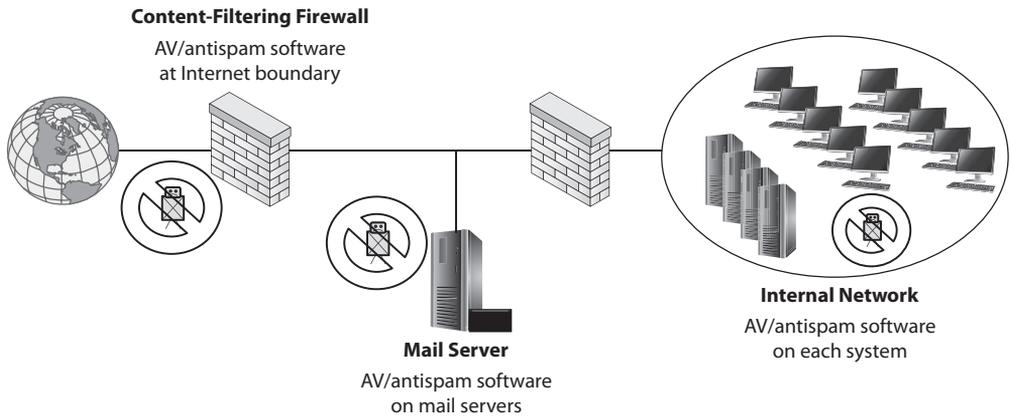


Figure 6-6 Content-filtering appliances

in almost no time, you'll be enjoying a hot meal. You don't need to understand how the microwave works to enjoy its benefits. Similarly, an administrator doesn't have to understand the complexity of the content-filtering appliance to use it and enjoy its benefits.

Most content-filtering appliances include additional features and capabilities that administrators can adjust. For example, if the organization uses it as a proxy server, it checks the validity of website certificates and blocks sites that fail these checks. It can also filter URLs, controlling what websites users can visit.



NOTE Chapter 4 presented proxy servers and how they can be used to control what websites users access. Proxy servers can also reduce Internet bandwidth usage.

Figure 6-6 also shows a defense in depth strategy to protect against malware by installing AV/antispam software at the boundary, on a mail server, and on all internal clients. Installing AV software in so many places might seem like overkill. After all, if most malware reaches systems via the Internet and e-mail, the content-filtering firewall and antivirus software on the mail server should catch it. Ideally, this is true. However, users can also inadvertently install malware via USB drives just by inserting the USB drive into their system. Additionally, zero-day exploits might get past the content-filtering firewall and/or the mail server. AV software on the systems might be able to detect this malware using heuristics-based detection methods. Together, these methods provide a defense in depth solution to protect against malware.



EXAM TIP Attackers commonly deliver malware over the Internet. Monitoring and scanning all Internet traffic is an important part of a defense in depth strategy to protect systems from malware.

Keeping Operating Systems Up to Date

It's also important to keep operating systems up to date. Malware often exploits software bugs or vulnerabilities in a system, and as long as the bug remains, the system is vulnerable. On the other hand, updated systems are not vulnerable to these bugs.

Vendors regularly write, test, and release patches for known vulnerabilities. For example, Microsoft releases patches on the second Tuesday of every month. The goal is to allow administrators time to plan for testing and installation of these patches on a regular basis. Although Microsoft does occasionally release patches for emerging threats outside of this schedule (out-of-band releases), these releases are relatively rare.

Chapter 9 talks about hardening systems, or making them more secure from their default configuration. Keeping systems up to date is a key part of hardening a system.

Scanners

Chapter 8 covers vulnerability scanners in more depth, but in short, vulnerability scanners can check systems for vulnerabilities and report their findings. Some common vulnerability scanners used by security professionals include Nessus, Nmap, and Saint.

Vulnerability scanners can check systems to verify that the operating system is up to date, has up-to-date antivirus software, and has an enabled firewall. A vulnerability scanner can also verify that the systems are not running certain services and protocols, depending on how the administrator configures the scanner.

Beware of Shortened Links

It used to be that you could easily spot malicious links in e-mails. For example, if you saw a link claiming to be from www.mcgraw-hill.com, you could hover over it to see the actual link was www.hackers.ru. In this case, clicking the link would not take you to [mcgraw-hill.com](http://www.mcgraw-hill.com) but instead would take you to a Russian website (because the *ru* top-level domain is in Russia).

As an example, Figure 6-7 shows an HTML-based web mail with the user hovering over the second link. Although the link appears to be www.mcgraw-hill.com, if you look in the status bar at the bottom of the browser, you can see the actual link is <http://www.hackers.ru>. Many e-mail programs allow the user to hover over the e-mail in the same way and see the actual link as a tool tip.

To prevent users from seeing the link, attackers sometimes use shortened URLs to mask the actual destination, as shown in the third link in Figure 6-7. Legitimate users also shorten very long URLs with tools such as bit.ly or tinyurl.com. You can go to the bit.ly website, plug in the long URL, and it gives you a short one. Now, when a user clicks the bit.ly link, it takes the user to the actual long link. Although this is useful (especially for 140-character tweets), attackers have used shortened links for attacks. End users don't know where this link actually takes them until they click the link. By that point, it may be too late.

Some tools are available to lengthen URLs. For example, LongURL (<http://longurl.org>) allows you to plug in a shortened URL and it expands it to a long URL. Users can copy

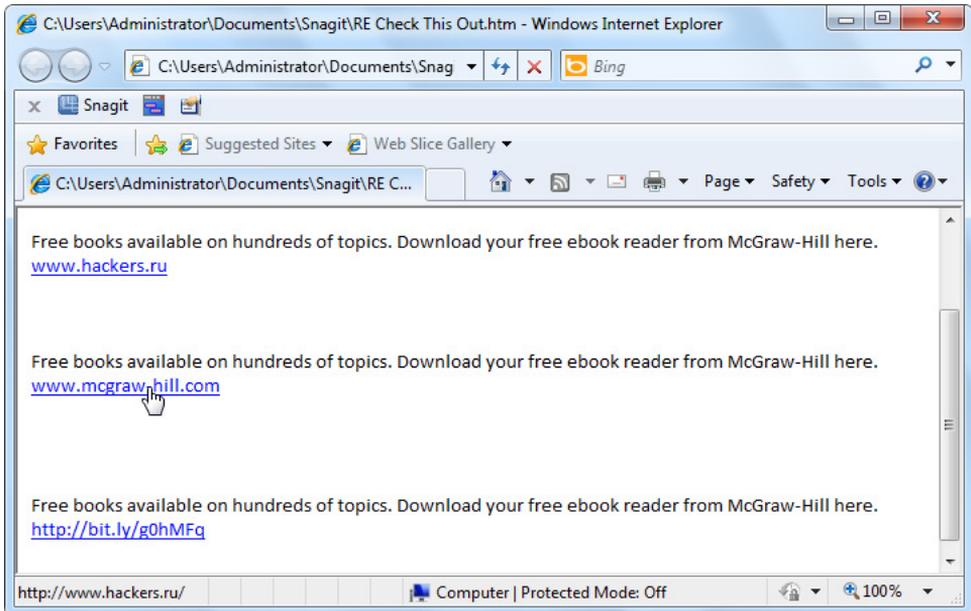


Figure 6-7 Embedding links in an e-mail

the shortened URL from their e-mail or from a tweet and paste it into the text box on the LongURL page. Figure 6-8 shows an example of lengthening a shortened URL.

In Figure 6-8, you can see that the bit.ly website shortened the `http://mcgraw-hill.com` URL to `http://bit.ly/dApKTQ`. Users have no idea where `http://bit.ly/dApKTQ` will take them. However, by pasting `http://bit.ly/dApKTQ` into the LongURL site, users can see the actual destination URL is `http://mcgraw-hill.com`.

Sandboxing

As mentioned previously, heuristic-based antivirus software sometimes uses a sandboxing technique. It runs new or untested code within an isolated environment. If the code is malicious, the code is unable to infect other applications or other areas of the system. Real-time scanning used by AV software uses sandboxing techniques.

Application sandboxing limits the environment where code can execute. One method that is often used is virtualization. It's possible to create virtual machines that are completely isolated from the host or the host's network. AV researchers often use virtual machines to analyze malware.

Least Privilege

Most operating systems require a user to have elevated permissions to install applications. Because of this, you can reduce the ability of malware to run on your system by logging on with an account that does not have elevated permissions. This follows a best practice

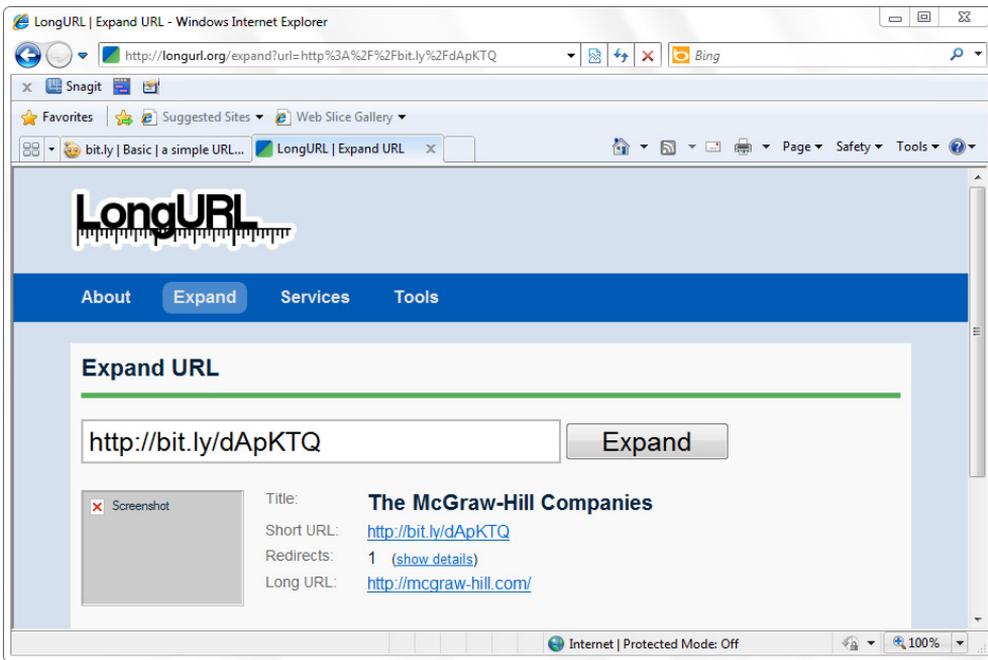


Figure 6-8 Expanding shortened links with LongURL

of giving users only the rights and privileges they need to perform their job—the principle of least privilege. If users don't install applications as a part of their job requirements, they should not have the privileges to do so, and this reduces the ability of the user to accidentally install malware.

A best practice that many organizations follow is to require administrators to have two accounts. One account is a regular user account with little or no privileges, and the second account has elevated administrator permissions. The administrator uses the regular user account for normal, day-to-day activities and uses the elevated user account only when performing administrative work. These steps reduce the opportunities for malware to be installed.



TIP Chapter 1 introduces the principle of least privilege as one of several fundamental security practices.

Software Security

Chapter 5 discussed several software security methods that can also help protect against malware. This includes using input validation techniques and application reviews for web applications that an organization hosts. It also includes verifying that risky code has been

signed with a digital certificate. Code signing helps reduce malware infections because criminals won't use code signing techniques to validate malware. It also helps identify modified code because a digital certificate won't identify code modified by a virus.

Application Whitelisting and Blacklisting

Another malware protection method is application whitelisting or blacklisting. Application whitelisting identifies specific applications that can run on a system and blocks all other applications. This can be effective in an organization that has a limited number of applications on systems, but it does make it more difficult to add new applications. Application blacklisting identifies specific applications that cannot run on a system. This becomes a little more difficult because it requires you to know all potential malicious applications. Chapter 9 covers application whitelisting and blacklisting as a method to secure endpoint devices.

Participating in Security Awareness and Training

Many users don't understand the risks associated with malware or how criminals try to trick them into installing malware. Because of this, it's important to educate users on safe computing habits. Security policies often mandate security awareness and training efforts to prevent security incidents before they occur. In other words, training is a preventive control. In contrast, detective controls attempt to detect incidents after they occur, and corrective controls attempt to reverse the effects of incidents after they occur.



TIP Chapter 5 mentioned the importance of educating users about social engineering. You can combine social engineering training with basic safe computing training.

The following list identifies some common items to repeat to users:

- Use spam filters and AV software.
- Don't open any attachments from unsolicited e-mails.
- Don't click any links in unsolicited e-mails.
- Be wary of all attachments and links in e-mails. Attackers often try to make e-mails look like they are from someone you know.
- Be wary of any warnings from a website, such as a warning that your system is infected with malware.
- Download files directly from the source website (not from a link provided through a third-party website).

Some organizations have security policies that require additional steps. For example, if a user notices a system is acting suspiciously as though it's infected with malware, the security policy may require that the user remove the computer from the network by disconnecting the network cable.

Common Vulnerabilities and Exposures

The Common Vulnerabilities and Exposures (CVE) list is a standardized list of known security vulnerabilities, exploits, and malware. Moreover, the CVE provides a standard way of naming information security vulnerabilities. The MITRE Corporation maintains the list, and the National Cyber Security Division of the U.S. Department of Homeland Security sponsors the CVE. You can access the CVE here: <https://cve.mitre.org/cve/>.



NOTE MITRE looks like an acronym, but it's actually just the name of the company, and is presented in all uppercase. Although some of the early employees previously worked at the Massachusetts Institute of Technology (MIT) providing research and engineering, MITRE is not a part of MIT.

The CVE has definitions for vulnerability and exposure that are worth knowing. A *vulnerability* is “a mistake in software that can be directly used by a hacker to gain access to a system or network.” An *exposure* is “a system configuration issue or a mistake in software that allows access to information or capabilities that can be used by a hacker as a stepping-stone into a system or network.”

Chapter Review

A significant security risk to computers and IT networks is in the form of malicious code, commonly called malware. Malware includes viruses, worms, Trojan horses, scareware, ransomware, keyloggers, logic bombs, rootkits, mobile code, backdoors, trapdoors, spyware, RATs, and malware hoaxes.

Viruses have a replication component to spread themselves and an activation component to deliver the payload of the virus. They use a variety of methods to thwart AV software and researchers. Stealth viruses attempt to hide themselves to escape detection. Armored viruses use techniques such as complex code and encryption to prevent a researcher from reverse engineering the code and discovering what the virus is doing. Polymorphic and metamorphic viruses change the virus file to escape detection through signature files.

Although viruses must execute, a worm can travel through a network without any human interaction. A Trojan horse appears to be one thing, such as a useful program, but is actually something different and malicious. Scareware is a type of Trojan. Scareware appears to be a useful antivirus program, but instead attempts to extort money from users. Ransomware locks a user's computer or data and demands ransom to return the control back to the user.

A keylogger captures keystrokes from users and can send the captured keystrokes to an attacker. Logic bombs execute in response to an event, such as a day or time. Rootkits take control of the operating system (at the kernel or root level) and have the ability to take complete control over what the user sees and does. A backdoor (or trapdoor) is code embedded into an application or operating system that allows an attacker to access data or even gain remote control access to the infected computer. Malware often installs RATs as backdoors into systems. Spyware is any software that is installed without the user's knowledge or consent and can be malicious. Most malware is delivered through the Internet via drive-by downloads, malvertising, and malicious e-mails.

The primary method of protecting systems from malware is the use of up-to-date AV software. AV software should be installed on every computer, and it's also highly recommended to install AV software at the boundary between the Internet and the internal network, and on e-mail servers. AV software uses signature definitions and heuristics to identify malware. Signature definitions identify known viruses based on a specific pattern within the virus. Heuristics attempt to identify malicious behavior of previously unknown viruses. Because viruses can mutate into different variations, it's important to keep virus definitions up to date. Additionally, it's important to keep an operating system up to date.

Spam filters and content-filtering appliances help filter out malware before it reaches the internal network. Content-filtering appliances can also filter other traffic, such as using URL filters to control which websites users can access, and performing firewall and intrusion detection services.

By educating users on the basics of safe computing habits, you can help them avoid risky computing behaviors that cause malware infections. This includes not opening attachments or clicking links from unsolicited e-mails, and being wary of all attachments and links in e-mails.

The MITRE Corporation maintains the CVE list, which provides a standardized method of describing malware. The National Cyber Security Division of the U.S. Department of Homeland Security sponsors the CVE.

Questions

1. What type of virus attempts to protect itself from reverse engineering and prevent antivirus researchers from analyzing the malware?
 - A. Armored virus
 - B. Polymorphic virus
 - C. Metamorphic virus
 - D. Multipartite virus
2. Which of the following malware types alters its own code to avoid detection by antivirus software?
 - A. Armored virus
 - B. Metamorphic virus
 - C. Polymorphic virus
 - D. Ransomware
3. What type of malware can spread without any user intervention?
 - A. Virus
 - B. Trojan horse
 - C. Worm
 - D. Spyware

4. A software application appears to have a useful purpose, but it includes malicious code. What does this describe?
 - A. A virus
 - B. A backdoor
 - C. A worm
 - D. A Trojan horse
5. After visiting a website, a user sees a pop-up indicating a virus has infected his system and offering free antivirus software. He downloads the free antivirus software, but finds that it won't clean the virus unless he purchases the full version. What does this describe?
 - A. Shareware
 - B. Rootkit
 - C. Freeware
 - D. Scareware
6. When Sally turns her computer on, she sees a screen indicating software has encrypted all of her data files. A message indicates she must pay \$300 within 48 hours to access the decryption key. What does this describe?
 - A. Logic bomb
 - B. Ransomware
 - C. Worm
 - D. Spyware
7. An employee configured malicious code to execute at midnight on February 2. What does this describe?
 - A. Logic bomb
 - B. Groundhog Day virus
 - C. Worm
 - D. Ransomware
8. What type of malware takes control of the operating system at the kernel level?
 - A. Trojan horse
 - B. Worm
 - C. Keylogger
 - D. Rootkit

9. A website developer wants to provide assurances to users that ActiveX controls used on the site are not malicious. What can provide this assurance?
 - A. Input validation
 - B. Code signing
 - C. Code review
 - D. Enabling cross-site scripting
10. What does antivirus software use to detect previously unknown viruses?
 - A. Signatures
 - B. Polymorphism
 - C. Heuristics
 - D. Armor
11. A virus is detected on a system based on the virus's behavior. What detected the virus?
 - A. Heuristics
 - B. A virus fingerprint
 - C. A virus filter
 - D. A signature
12. What should users do to ensure that antivirus software can detect recently released viruses?
 - A. Update signatures
 - B. Update the operating system
 - C. Update the AV software
 - D. Regularly purchase new AV software
13. Of the following choices, how is malware most often delivered today?
 - A. Over the Internet
 - B. Via an intranet
 - C. Via USB drives
 - D. Through company policies
14. A company authorizes users to transport data from work to home using USB drives. What's the best method of protecting systems from malware without affecting the user?
 - A. Install AV software on the network firewall
 - B. Install AV software on the e-mail server
 - C. Install AV software on each user's work computer
 - D. Prevent users from using USB drives

15. Of the following choices, what is the best technique you can implement on an e-mail server to reduce infection through e-mail?
 - A. Block all e-mail
 - B. Add a spam filter
 - C. Add a polymorphic filter
 - D. Remove all attachments
16. Of the following choices, which one is NOT a valid method to reduce malware infections?
 - A. Don't open attachments from unsolicited e-mails.
 - B. Don't click links in unsolicited e-mails.
 - C. Don't send encrypted personal information via e-mail.
 - D. Don't follow shortened links from unknown sources.
17. Of the following choices, which one is a principle that prevents users from accidentally installing malicious software on their systems?
 - A. Nonrepudiation
 - B. Least privilege
 - C. Separation of duties
 - D. Accountability
18. Of the following choices, what network device can filter e-mail, spam, and malware?
 - A. Packet-filtering firewall
 - B. Proxy server
 - C. An intrusion detection system
 - D. Content-filtering appliance
19. Your organization mandates security training for users within its security policy to educate users about malware and methods to prevent malware infections. What is the best description of this effort?
 - A. A detective control
 - B. A corrective control
 - C. A preventive control
 - D. A technical control
20. What provides a standardized method of describing malware?
 - A. The Consortium of Antivirus Vendors (CAV)
 - B. The Consortium of Virus Authors (CVA)
 - C. The National Institute of Standards and Technology (NIST)
 - D. The Common Vulnerabilities and Exposures (CVE) list

Answers

- 1. A.** An armored virus attempts to prevent an AV researcher from reverse engineering it to determine what it is doing and how it is doing it. Although polymorphism and metamorphism can make it harder to reverse engineer a virus, they aren't the best answer, because these techniques primarily make it harder for AV software to detect the virus. A multipartite virus uses multiple methods of attack.
- 2. B.** A metamorphic virus changes or mutates its code as it replicates itself to prevent detection. An armored virus uses techniques such as encryption to make it more difficult for AV researchers to decompile the virus. A polymorphic virus changes the file, but not the code. Ransomware takes over a user's computer and demands a monetary ransom to return control back to the user.
- 3. C.** Worms spread through a network without any user intervention. Viruses, Trojan horses, and spyware all require some level of interaction.
- 4. D.** A Trojan horse appears to be something useful to the user but includes malicious code or malware. While Trojans often include viruses and backdoors, not all viruses and backdoors come from Trojans. Worms travel over the network and are not embedded in software applications.
- 5. D.** Scareware is malware that scares users into thinking a virus has infected their system and encourages them to install a free download. The free download appears as antivirus software that doesn't remove viruses unless users pay, but it often includes malware itself. Shareware is software that users are free to try and pay for if they like it and continue to use it. A rootkit takes over the system with root-level privileges. Freeware is free software.
- 6. B.** Ransomware takes control of a user's computer or data and demands a ransom to return control to the user. This scenario describes CryptoLocker. A logic bomb is malware that executes in response to an event such as a specific date and time. Worms infect computers over a network, and while worms deliver malware, not all worms include ransomware. Spyware is software installed on a user's system without the user's knowledge with the goal of spying on the user, not extorting money from the user.
- 7. A.** A logic bomb is malware that executes in response to an event such as a specific date and time. While February 2 is Groundhog Day, the scenario doesn't describe a Groundhog Day virus. Worms infect computers over a network, not on a specific day. Ransomware takes control of a user's computer or data and demands a ransom from the user.
- 8. D.** A rootkit is a set of programs that runs on a system, largely undetected, because it runs at the kernel level or root level of the operating system. A Trojan horse is malware that looks like one thing but is something else. A worm is a type of malware that spreads through a network without any user intervention. A keylogger captures keystrokes from users.

9. **B.** Code signing digitally signs ActiveX controls and provides assurances to users of who created the control and that it hasn't been modified. Input validation helps prevent injection attacks, but it's used to protect the website, not provide assurance to users. Code review is a valuable tool to detect problems with applications before an organization releases them. Cross-site scripting is an attack and would not be enabled.
10. **C.** Antivirus software uses heuristics to detect previously unknown viruses. Signatures detect known viruses. Polymorphism and armor are techniques used by virus authors to prevent the detection of a virus.
11. **A.** Heuristics can detect malware based on the behavior of the malware and are designed to detect previously unknown viruses. There's no such thing as a virus filter or a virus fingerprint, although a virus signature does uniquely identify known malware similar to how a fingerprint can identify a person.
12. **A.** Antivirus software uses signature definition files to detect viruses, and these signatures must be regularly updated. It's not necessary to update the operating system, update the AV software, or purchase new AV software to detect recently released viruses.
13. **A.** The common way attackers deliver malware is over the Internet. While some attacks can come from internal intranet sources, they do not compete with the volume of attacks from the Internet. Unsuspecting users transmit viruses with USB drives, but this isn't as common as virus delivery over the Internet. Company policies would not deliver viruses.
14. **C.** Installing AV software on each user's work computer provides the best protection against a user inadvertently transporting malware from home to work. Installing software on the network firewall and on an e-mail server is a good practice, but it won't help if the virus is transported via a USB drive. Preventing the users from using USB drives will affect the users.
15. **B.** The majority of malware comes through spam, so a spam filter can reduce infections through e-mail. An e-mail server isn't very useful if it blocks all e-mail or removes all attachments. E-mail servers don't have polymorphic filters.
16. **C.** If you need to send personal information via e-mail, the best choice is to send it in an encrypted format. All of the other choices are valid methods to reduce malware infections.
17. **B.** The principle of least privilege specifies that users are given rights and privileges to do their job but no more. If a user doesn't need to install applications, the user is not given permission to do so, which reduces the possibility of the user accidentally installing malware. Nonrepudiation prevents a person from denying an action. Separation of duties divides tasks so that no single person or entity controls an entire process. Accountability ensures that user actions can be tracked and monitored.

18. **D.** Content-filtering appliances can filter e-mail to remove spam and malicious attachments. They can also act as a proxy server or an intrusion detection system (IDS). A packet-filtering firewall can only examine packets. A proxy server can filter websites but not e-mail. An IDS detects attacks, but not e-mail and spam.
19. **C.** Training is a preventive control because it attempts to prevent incidents from occurring. Detective controls attempt to detect incidents, and corrective controls attempt to reverse the effects of an incident. Technical controls use technology to implement the control, but training doesn't require technology.
20. **D.** The CVE is maintained by the MITRE Corporation and provides a standardized method of describing security vulnerabilities, exploits, and malware. There is no such thing as the Consortium of Antivirus Vendors (CAV) or Consortium of Virus Authors (CVA). NIST is a U.S. government entity that regularly publishes standard publications related to IT security, standards, and practices, but it does not maintain the CVE.

Risk, Response, and Recovery

In this chapter, you will learn about

- The definition of risk, threat sources, threat events, and vulnerabilities
- Primary methods of managing risk
- The definition of residual risk
- Risk assessment steps
- Differences between quantitative and qualitative analysis
- Steps in incident response

Defining Risk

The simple definition of *risk* is the probability or likelihood that a threat will exploit a vulnerability and cause a loss. A *threat* is any activity that can be a possible danger. A *vulnerability* is a weakness, and a *loss* represents a negative event for an organization. Figure 7-1 shows these components and their relationships with each other.

A common formula used to represent risk is

$$\text{Risk} = \text{Threat} \times \text{Vulnerability}$$

This formula isn't a mathematical algorithm. However, it does show the relationship between threats and vulnerabilities and risk. Small threats and insignificant vulnerabilities result in a low level of risk. More severe threats and serious vulnerabilities result in a higher level of risk.

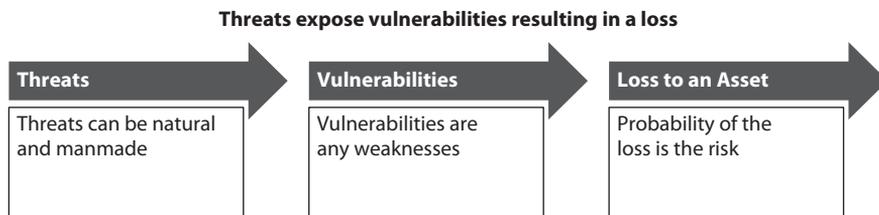


Figure 7-1 Components of risk

When the threat and the vulnerability are combined (a threat exploits a vulnerability), it results in a loss. More significant threats or vulnerabilities can easily result in higher losses from the risks. In other words, some risks have higher impacts. However, if you can reduce the vulnerabilities, you reduce risk, resulting in lower potential losses.



TIP It's rarely possible to reduce threats. You can't stop an attacker from attacking, but by reducing vulnerabilities, you reduce the attacker's chances of success. In other cases, you can reduce the impact of the risk to limit losses.

For example, malware represents a threat to any computer. If a computer does not have antivirus (AV) protection, it is vulnerable to malware threats. If a virus infects an unprotected system, it can modify, destroy, or steal data, representing a loss to the owner of the computer. An organization can mitigate, or reduce, the vulnerability by using AV software and keeping it up to date with current antivirus signatures.



EXAM TIP Risk is the probability or likelihood that a threat will exploit a vulnerability and cause a loss. Risk mitigation is the practice of reducing risk. The primary method of reducing risk is to reduce or eliminate vulnerabilities by implementing safeguards or controls. It's also possible to reduce risk by reducing the impact of the risk.

Total risk is the combined risk to all of the organization's assets, including all the threats and vulnerabilities. Total risk also includes the value of these assets.

NIST

The U.S. National Institute of Standards and Technology (NIST) includes a Computer Security Division and an Information Technology Laboratory (ITL) within that division. The ITL performs research to advance the development and productive use of information technology and regularly publishes documents known as Special Publications (SPs). The SP 800 series of reports includes the results of this research performed in collaboration with industry, government, and academic organizations.

You can review a full list of the NIST ITL SP 800 series documents here: <http://csrc.nist.gov/publications/PubsSPs.html>. You can also download any of these documents as PDF files.

NIST SP 800-30, *Guide for Conducting Risk Assessments*, provides a more comprehensive definition of risk. It defines risk as “a function of the likelihood of a threat event's occurrence and potential adverse impact should the event occur.” It also mentions that



Figure 7-2 Risk model

a threat source initiates a threat event, which exploits a vulnerability, causing an adverse impact. Figure 7-2 shows the risk model outlined in NIST SP 800-30, and the following list explains these elements:

- **Threat source** A threat source is anyone or anything that can initiate a threat event. As an example, an attacker is an adversarial threat source.
- **Threat event** A threat source initiates a threat event. For example, an attacker can launch a buffer overflow attack against a web server. The attacker is the adversarial threat source, and the buffer overflow attack is the threat event.
- **Vulnerability** The extent of the threat event is related to the extent of the vulnerability. For example, if a website doesn't use input validation techniques, it is more vulnerable to a buffer overflow attack than a website that does use input validation techniques. Security controls, such as input validation techniques, reduce vulnerabilities.
- **Adverse impact** An adverse impact affects the confidentiality, integrity, and/or availability of IT systems or data, and in some cases can affect the image and reputation of the organization.



TIP Chapter 1 presented the concepts of confidentiality, integrity, and availability (CIA). To recap, maintaining confidentiality means that unauthorized individuals don't have access to data or systems. Maintaining integrity means that unauthorized individuals can't modify data and systems. Maintaining availability means that data and systems are up and operational when needed.

While the SP 800-30 definitions and explanations are more technical, they do repeat the common themes of threats, vulnerabilities, impacts, and losses.

Identifying Threat Sources

While it's common to think of an attacker as a threat source, there are several other threat source types. NIST SP 800-30 categorizes them as adversarial, accidental, structural, and environmental:

- **Adversarial threat sources** Adversarial threat sources include individuals (internal or external) or groups that purposely attempt to exploit weaknesses. Internal threats are from malicious employees, contractors, or any other personnel granted access to an organization's resources, and have the capability of launching

an attack from within. External threats are from entities that don't have direct access to the organization's resources, such as attackers in a different city or country. Groups include both loose-knit groups of attackers and highly organized attackers (such as state-sponsored advanced persistent threats). Groups can also include competitors, suppliers, partners, and criminals.

- **Accidental threat sources** Accidental threat sources include erroneous actions taken by individuals in the course of executing their everyday responsibilities. For example, a user can ruin a laptop simply by spilling coffee onto the keyboard. Similarly, if someone accidentally kills the power or cuts communications cables to key systems, it results in a loss of availability. Users can accidentally cause damage to data, and the level of damage they cause is directly related to their level of access. Regular users have limited access to resources, so they can only cause a limited amount of damage (such as accidentally deleting some of their files). Privileged users such as administrators have much more access, so they can accidentally cause much more damage (such as deleting an entire customer database). While IT security topics often stress intentional threats, it's important to realize steps need to be taken to protect assets from accidental threats also.
- **Structural threat sources** Structural threat sources include failures of IT equipment, environmental controls, and software. As an example, if the environmental controls fail in a data center resulting in a loss of air conditioning, it can result in damaged equipment. Without air conditioning, servers get hot, and they can overheat and fail. Because of this, many organizations will power down servers when air conditioning fails. This protects the equipment, but also results in an immediate loss of availability.
- **Environmental threat sources** Environmental threat sources include natural and manmade disasters. Natural causes include weather events such as lightning storms, hurricanes, tornadoes, and floods, and can include non-weather events such as earthquakes or tsunamis. One example of a manmade disaster is a fire. Environmental threat sources also include failure of external services such as a telecommunications link, or power outages.



TIP Appendix D in SP 800-30 includes additional information on threat sources, including several tables listing threat sources and sample assessment scales used to categorize them.

Identifying Threat Events

Threat events include any type of activity or event that can result in a loss of confidentiality, integrity, or availability to a system. SP 800-30 matches threat events with threat sources to identify threat scenarios. A threat scenario is one or more threat events, caused by one or more threat sources, that cause an adverse impact.

Many times, you'll see the threat source and threat event combined as simply a threat. However, when performing risk assessments, security professionals often separate the two

so that they can more accurately identify who the threat is (the threat source) and what they might do (the threat event).

NIST SP 800-30 includes an extensive list of potential adversarial threat events, grouped into several categories. These categories focus on known attack methods and activities, and the following list includes some of them:

- **Perform reconnaissance and gather information** Adversaries use a variety of techniques to learn more about their target. These techniques include gathering information from public sources and using tools such as sniffers and scanners to gather technical information. Attackers often try to trick users into installing malware on their computers. When successful, it allows the attacker to remotely control the computer and perform reconnaissance on the infected computer's internal network. These attacks from within the organization become more difficult to track because most detection tools focus on attacks coming from outside the organization.
- **Craft or create attack tools** Attackers create a wide variety of different attack tools. For example, they often use phishing attacks (including spear phishing and whaling attacks) to collect information and trick users into installing malware on their systems. They have been quite creative in their methods, including spoofing e-mails to appear to originate from funeral homes, delivery services, and more. Their goal is to get the user's interest. To complement these methods, attackers develop malware, which they include as attachments.
- **Deliver/insert/install malicious capabilities** Once attackers create the malware, they find ways to get the malware installed on user systems. Attackers often take the time to analyze methods that have worked before and create similar methods. For example, many users don't recognize the risk of opening a ZIP file, so attackers often embed malware within ZIP files, which they send as attachments with phishing e-mails. They have also attacked websites and installed malware on them. The malware then attempts to download itself automatically when users visit.
- **Exploit and compromise** Exploits include both technical and nontechnical methods. For example, social engineers circumvent physical security methods by tailgating closely behind authorized individuals to access secure locations. Technical methods exploit known vulnerabilities in unpatched systems. Additionally, attackers are always on the alert for zero day vulnerabilities that don't have patches. Many exploits allow attackers to assume control of compromised systems from remote locations.
- **Achieve results (such as causing adverse impacts and obtaining information)** Once attackers successfully exploit systems, they go after their target. This is often to obtain sensitive information from the victims, cause damage to the systems, and/or infiltrate a network to steal data. Adverse impacts include any results that affect the confidentiality, integrity, or availability of systems and data.

- **Maintain a presence or set of capabilities** Once attackers are within a system or network, they take steps to obscure their actions. This includes using methods to bypass or circumvent intrusion detection systems and auditing capabilities. Security firm Mandiant has documented several cases where an advanced persistent threat (APT) maintained a presence within networks for several months and sometimes for several years.

It's worth stressing that the previous list does not include all threat events. Attackers have a wide variety of methods they use to attack organizations. As security professionals find ways to thwart these attacks, attackers modify their methods using new and modified attacks.

Understanding Vulnerabilities

Vulnerabilities are any weaknesses in a system, network, infrastructure, or organization. A vulnerability can also be the absence of a security control or countermeasure. For example, a firewall protects systems and networks from Internet threats. A network without a firewall indicates a vulnerability. NIST SP 800-30 defines a vulnerability as follows:

A vulnerability is a weakness in an information system, system security procedures, internal controls, or implementation that could be exploited by a threat source.

If a threat can exploit these weaknesses, it results in a loss. Some examples of vulnerabilities include the following:

- **Poor or nonexistent AV protection** Computer systems need antivirus software with up-to-date virus signatures. If systems don't have AV protection, malware can infect them, resulting in costly losses.
- **Terminated employees** If employee accounts are not disabled or deleted after an employee is terminated, the account can be used to attack a system. The employee may be able to access the system and cause damage, or another employee may use the account to perform actions. If another employee accesses the account, logs will inaccurately attribute the actions to the terminated employee.
- **Weak access controls** Access controls include identification, authentication, and authorization techniques. Weak access controls allow unauthorized individuals to access, modify, and/or destroy data.
- **Poor or nonexistent change-management practices** An unintentional side effect of making a change to a system is that the change causes one or more problems. Strong change-management practices examine prospective changes to predict the side effects and can mitigate problems caused by changes.
- **Poor hardening practices** Systems should be hardened, or made more secure from their default configuration. This includes changing defaults such as default passwords, implementing firewalls, keeping operating systems and applications up to date with patches, and more.

- **Lack of redundancies for key systems** Any single point of failure can result in an outage of an entire system if it fails. Administrators can add redundant technologies for disk drives, servers, communication lines, power, air conditioning, and even entire sites or locations. Adding redundancies allows systems to tolerate faults. For example, if a single drive in a RAID-6 array fails, the disk subsystem can tolerate the fault and continue to operate.
- **Uneducated users** Social engineers can trick uneducated users into giving up key information. On the other hand, social engineers aren't as successful when they try to trick educated users. Similarly, educated users are less likely to be tricked into installing malware onto their system by opening attached files or clicking links in malicious e-mail.



TIP The preceding list is by no means complete, but it does provide some common examples of vulnerabilities. An important point to remember about any vulnerability is that you can reduce weaknesses by implementing controls or safeguards. By reducing the vulnerabilities, you reduce the possibility of the risk resulting in a loss.

It's worth mentioning that an organization doesn't necessarily have to eliminate all vulnerabilities. For example, some legacy systems might be critical to the business, but include known vulnerabilities that cannot be mitigated. Similarly, some patches adversely affect applications. While testing a patch, administrators might detect the problem and report it to management, and management might decide to accept the risk of leaving the system unpatched.

Threats often exploit the weakest link in your organization, and once a threat exploits a vulnerability, the damage is done. A key method used to minimize losses is to identify and reduce vulnerabilities.

Understanding Impact

The probability that a threat can exploit a vulnerability provides the likelihood of occurrence. However, it's also important to examine the potential impact of the risk occurring. The *impact* refers to the magnitude of harm that could result if a threat exploits a vulnerability. You can identify the magnitude of harm by estimating the potential loss of confidentiality, integrity, or availability. In some cases, the impact is negligible, making the risk a low priority. In other cases, the impact is quite high, making the risk a very high priority.

Consider the risk of a fire. Overall, if personnel within an organization use safe practices, the risk of a fire is low. Fires do occur, so organizations often install fire detection and fire extinguishing equipment to automatically detect and extinguish a fire. Still, if these methods fail, a fire can destroy most of the contents of a building in a very short time. The impact can be catastrophic. In this case, the likelihood of occurrence for a fire is low, but the potential impact is high, so organizations purchase insurance to reduce the overall risk by reducing the impact.

In contrast, consider the likelihood that students in a classroom might steal a \$10 mouse. The mice might be easy to remove and hide in a book bag, so the likelihood might be high. However, the impact of the loss is very low, so the overall risk is low.

Risk assessments consider the impact of the risk, although they use different methods such as quantitative and qualitative risk analysis methods. A quantitative risk assessment uses costs and monetary values to identify the impact. A qualitative risk assessment uses subjective opinions of experts to identify the impact. Both are covered later in this chapter.

Managing Risk

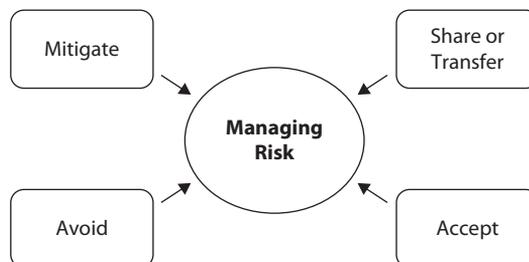
Risk management attempts to identify, assess, and mitigate risk to an acceptable level for the organization. As previously discussed, two key components of risk are threats and vulnerabilities, so these are key components to address when attempting to manage risk. It's important to realize that risk management is an ongoing process. Threats constantly change and evolve, and risk management practices need to change and evolve with them.

It isn't possible to reduce the overwhelming majority of threats. However, you can reduce vulnerabilities and reduce the impact of threats. For example, you can't stop a hurricane or an earthquake. However, you can reduce the impact of these natural threats by taking additional precautions, such as building stronger buildings or simply building them outside of a threat zone. Similarly, you can't stop an attacker from trying to launch attacks against your organization, but you can take steps to reduce the attacker's possibilities of success by strengthening security practices and reducing vulnerabilities.

Overall, an organization can select from several choices in response to different risks. Figure 7-3 shows these choices, and the following list describes them:

- **Mitigate** You mitigate a risk by reducing vulnerabilities, and you reduce vulnerabilities by implementing controls or safeguards. For example, you can add AV software to reduce the risk of malware infection. You can add intrusion detection systems to detect and prevent network attacks from penetrating your network.
- **Avoid** You can avoid a risk by simply avoiding the activity that introduces the risk. For example, if you want to avoid the risks of a hurricane, you can move your location to an inland location. If you want to avoid risks associated with hosting a website on the Internet, you can choose not to host a website.

Figure 7-3
Managing risk



- **Share or transfer** You can share a risk with another party, or transfer the risk to another party. The most common method of transferring a risk is with insurance. For example, you can purchase fire insurance, and if a fire destroys a building or a business, the insurance company covers the loss. It isn't always possible to transfer all responsibility of the risk to another party. For example, buying insurance might help an organization cover some losses, but the insurance doesn't necessarily cover all the losses.



TIP Previous documentation listed *transfer* as a risk management method but did not include *share*. However, that has been changing in recent years. Many security professionals recognize that it often isn't possible to transfer the entire risk, but it is possible to transfer some of it, or share the risk with an outside entity. For this reason, this option has changed from *transfer* to *share or transfer*.

- **Accept** You can also choose to accept the risk and its potential losses. This is appropriate when the asset value is low or if the cost to reduce the risk is higher than the value of the asset. For example, someone may steal a mouse, but the cost of securing mice to prevent theft is higher than the cost of the mice. Additionally, it's appropriate to accept the remaining risk after taking steps to reduce the risk to an acceptable level.

Residual Risk

Residual risk is the amount of risk that remains after mitigating risk. In other words, total risk – mitigated risk = residual risk, as shown in Figure 7-4. It's important to ensure that residual risk is at an acceptable level to the organization. If the residual risk is too high, then the organization needs to find additional security controls to reduce the risk.

An old saying is that the only way to keep a computer 100 percent safe is to never plug it in and turn it on. Of course, the computer isn't very useful in this state. Instead, management decides what controls to implement to reduce the risk to an acceptable level. These decisions result in mitigating some risks and accepting other risks.

With this in mind, it's important to remember that management is responsible for any residual risk. Additionally, management is responsible for any losses associated with



Figure 7-4 Residual risk

residual risk. Management chooses which mitigation methods to implement and which methods to reject. Their goal is to reduce risk to an acceptable level based on the needs of the organization.



EXAM TIP Residual risk is the risk that remains after reducing risk to an acceptable level. Management is responsible for approving (or rejecting) controls that can mitigate risk. Management is also responsible for any losses resulting from residual risk.

Of course, different organizations have different risk levels. A company that takes anglers out fishing won't intentionally go into a war zone. However, a naval warship will intentionally go into harm's way when necessary, even though the crew will take steps to reduce risk.

Management has an overall view of the entire organization, including understanding the organization's mission and vision. Security personnel won't necessarily know all the reasons why management might decide to implement one control to reduce a risk but reject another control and accept a different risk. Still, the key is that management is responsible for making these decisions, and is responsible for any losses resulting from residual risk.

Identifying Assets

An important first step in risk management is identifying an organization's valuable assets. Assets with high value to the organization need to be protected by implementing controls or safeguards to reduce risks. While this seems obvious, it's not always apparent to security professionals what assets the organization considers to be of the highest value. For example, a web server that looks just like any other server in the server room may be quietly generating thousands of dollars an hour in revenue. Technicians working in the server room might not recognize the server's value and instead prioritize other servers during a security incident.

Assets are valued using both tangible and intangible values. For example, consider a company that sells products via a website and generates an average of \$6,000 an hour. If the website goes down for an hour, the company loses \$6,000 in direct, tangible revenue. However, many other intangible losses aren't so apparent. Customers who visited the website while it was down may choose to shop at another company. These lost customers represent intangible losses. This includes the money expended to attract these customers and the loss of future sales to these customers. It also includes the money required to gain new customers. If a company suffers a sustained or highly public loss, it needs to expend advertising dollars to regain customer goodwill. Otherwise, even customers unaffected by the outage may choose to go elsewhere.

Of course, there's no need to waste time evaluating risk if an asset has very little value to an organization. It's relatively inexpensive to replace low-value assets if necessary, so they don't require much time or energy to protect.

Risk Visibility and Reporting

A key element of risk management is risk visibility and reporting. Ideally, an organization will be aware of the risks that threaten the organization's resources, but it does take time and effort to identify those risks. Processes such as risk assessments help identify these risks so that the risks are known and visible. Continuous monitoring systems such as antivirus software, data loss prevention (DLP) systems, firewalls, and IDSs monitor the resources and create reports or alerts when they detect problems.

Risk Register

A risk register (or risk log) is a central depository for known organizational risks and is a key method used to share threat intelligence. The risk register includes details about the known risks and can consolidate data from multiple sources. Table 7-1 shows a sample table of a risk register for a web server hosting an e-commerce site. It shows sample threats along with estimated values for the probability of each of the threats occurring and the impact if it does occur. The risk score is the product of the probability and impact columns and helps security professionals prioritize the risk.



NOTE Chapter 5 covers various attacks, including SQL injection and buffer overflow attacks. It also includes methods to mitigate these attacks, such as input validation and keeping systems up to date with current patches.

Table 7-1 quickly shows that the greatest risk of those listed is a buffer overflow attack, with a score of 56. The table also shows what controls are in place. Because input validation is being used to prevent a SQL injection attack, the risk score is relatively low. If input validation methods weren't in place, the score would be much higher.

While Table 7-1 is limited by the width of the book's page, actual risk registers don't have the same limitations. It's easy to add additional columns, such as

- **Description** Full descriptions of the risks so that personnel using the risk register understand them.
- **Planned controls** These are any controls that management has approved but have not been implemented yet.

Name	Probability (P 1-10)	Impact (I 1-10)	Risk Score P × I	Controls	Remarks
SQL injection attack	1	8	8	Input validation	Control in place
Buffer overflow attack	7	8	56	Keep systems up to date	Systems not kept up to date
Data breach	2	9	18	Data on internal server	Control in place

Table 7-1 Sample Risk Register

- **Risk score with planned controls** Experts can predict the probability and impact of the risk with planned controls in place.
- **History of attacks** This column can provide links to an additional table that documents past attacks, along with explanations of how they were successful and what was implemented after the attacks.
- **Incident response** This provides details on the best response to the incident to mitigate losses.

An organization can create its own internal risk register, but there are also publicly available risk registers that are quite comprehensive. Chapter 6 mentioned the Common Vulnerabilities and Exposures (CVE) list. This is a standardized list of known exploits, security vulnerabilities, and malware organized as a massive dictionary. It provides standardization for naming vulnerabilities and exposures. The National Vulnerability Database (NVD) includes CVE data plus much more, such as checklists, alerts, and vulnerability notes. Both the CVE and NVD are types of risk registers.

Performing Risk Assessments

A risk assessment is a point-in-time evaluation. It looks at the current situation and attempts to determine what risks exist and how to address them. As an example, you can perform a risk assessment to determine risks for e-mail servers in your organization. This assessment could examine current threats, system configuration, interconnectivity, the amount of e-mail processed, fault-tolerance capabilities, power and air conditioning requirements, and more. However, after the assessment is completed, any of these elements can change, affecting the results of the assessment.

One of the benefits of performing risk assessments is that it allows security personnel to identify and prioritize risks. This helps ensure the organization dedicates risk management resources to the most serious risks. A risk assessment will use either a quantitative analysis or a qualitative analysis.



EXAM TIP A risk assessment looks at risks at a specific time. In contrast, risk management is an ongoing practice. Organizations repeat risk assessments periodically to ensure that current risks are assessed.

Quantitative Analysis

A quantitative analysis uses numerical-based data such as monetary figures to identify the actual cost associated with a risk. There are several variables associated with a quantitative analysis. If you can identify the values of these variables, prioritizing different risks becomes a simple matter of plugging the values into a formula. These values also help to justify costs for controls and countermeasures.

The following terms are important to understand when using a quantitative analysis:

- **Single loss expectancy (SLE)** This value identifies the actual loss of a single occurrence. For example, if a revenue-generating website is down for an hour, the company could determine that it results in a loss of \$10,000.

- **Exposure factor** The exposure factor identifies the magnitude of loss as a percentage of an asset's value. You might see this associated with the SLE. As an example, imagine that a building has a value of \$500,000 and insurance covers 75 percent of the replacement cost. The exposure factor is 25 percent (or \$125,000).



TIP The SLE would include both tangible and intangible costs. Tangible costs would include the actual lost revenue while the website is down. Intangible costs include the costs associated with lost customers or advertising to rebuild the company image.

- **Annual rate of occurrence (ARO)** This identifies how often the SLE is expected to occur in a given year based on the current environment. You can use historical data to determine the ARO. For example, if the website experienced three one-hour outages each year in the past two years, it will probably experience three one-hour outages this year if something isn't changed.
- **Annual loss expectancy (ALE)** You calculate the ALE by multiplying SLE \times ARO. For example, if the SLE is \$10,000 and the ARO is three, the ALE is $\$10,000 \times 3$, or \$30,000.
- **Cost of control** A control will mitigate or reduce a risk. For example, if faulty hard drives caused an outage, you can replace them with a hardware RAID-6. If the RAID-6 costs \$5,000, then the cost of the control is \$5,000. Some controls have ongoing costs to consider. For example, administrators might need additional training to install and manage some systems, such as dedicated firewalls or unified threat management devices.

You can compare the ALE with the cost of the control to determine if it makes fiscal sense to purchase the control. In general, if the cost of the control is less than the ALE, the cost is justified. If the cost of the control exceeds the ALE, it may not be justified. If the cost of the control is about the same as the ALE, then the organization can perform a detailed analysis to determine the actual return on investment (ROI) for the control.

Consider the example of website outages caused by hard drive failures. Experts might predict that adding the hardware RAID-6 will eliminate these outages. You can then use the following information to see if it makes sense to purchase the hardware RAID-6:

- ALE before control: \$30,000
- ALE after control: \$0
- Cost of control: \$5,000
- Savings (ALE before control – ALE after control – cost of control): \$25,000

If the company does nothing, it can expect to lose \$30,000. If the company implements the control, it eliminates the outages related to hard drives and reduces the ALE to \$0. In other words, if the company spends \$5,000, then it can expect to save \$30,000, for a net savings of \$25,000. The cost of the control is justified.

It's also possible that by implementing a control, you don't expect to eliminate the outages, but instead you expect to reduce the ARO. Using the previous example, experts might predict that adding a hardware RAID-6 will reduce the ARO from three to one. Instead of failing three times, the website might fail once due to a hard drive failure. An ARO of one results in an ALE of \$10,000 ($1 \times \$10,000$). Is the cost of the RAID-6 still justified?

- ALE before control: \$30,000
- ALE after control: \$10,000
- Cost of control: \$5,000
- Savings (ALE before control – ALE after control – cost of control): \$15,000

You've reduced the ARO from three to one, so you've reduced the ALE from \$30,000 to \$10,000. This provides an initial savings of \$20,000. The cost of the control is \$5,000, so this provides a net savings of \$15,000. If you did nothing, you'd lose \$30,000, so this cost is still justified even if it doesn't eliminate the risk.

On the other hand, if another risk has an ALE of \$2,000 and the cost of the control is \$22,000, the control is not justified. In this example, the control will reduce the ALE to \$0:

- ALE before control: \$2,000
- ALE after control: \$0
- Cost of control: \$22,000
- Savings (ALE before control – ALE after control – cost of control): negative \$20,000

You're spending \$22,000 on a control to prevent a \$2,000 loss. In this situation, management would typically decide not to implement the control, but instead accept the risk.



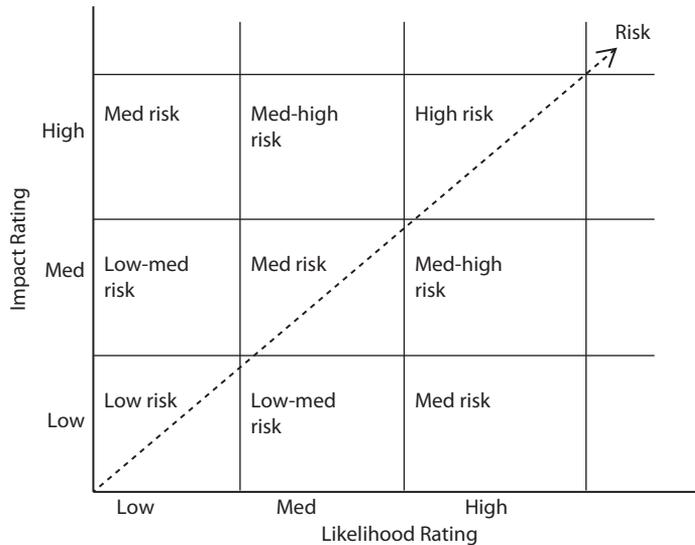
EXAM TIP When preparing for the SSCP exam, it's important to have a solid understanding of SLE, ARO, and ALE. You should also be able to evaluate the cost of a control to determine whether it is justified.

An advantage of the quantitative analysis is that it provides a clear metric in the form of a monetary figure. You can use this to analyze the actual value of the control. However, the numbers used to create a quantitative analysis are not always easy to obtain. This is especially true when the SLE includes intangibles.

Qualitative Analysis

A qualitative analysis is subjective and often simply categorizes a risk using words such as low, medium, and high. For example, Figure 7-5 shows the relationship between impact and likelihood ratings and how to use them to determine overall risk. Notice that there's a

Figure 7-5
Qualitative
analysis based
on impact and
likelihood ratings



direct relationship between likelihood and impact. As the likelihood and impact increase, the risk also increases.

In Figure 7-5, you can see that a low impact combined with a high likelihood results in a medium risk. It's important to realize this is just a guide and that the security professional doing the assessment will likely add in some subjective reasoning. As mentioned previously, a fire has a low likelihood but a high impact. However, it may be classified as an overall high risk because the result of a fire can be catastrophic.

It's also possible to assign values to the ratings. This allows you to prioritize the risks based on the result. For example, imagine that you've assigned a value of 1 for low, 5 for medium, and 10 for high. You can then use these values to determine a risk score by multiplying the likelihood rating by the impact rating. Table 7-2 shows a risk level matrix that combines these ratings and shows the resulting risk scores. Notice that data such as this can be included in a risk register (as shown in Table 7-1 earlier).



TIP An organization can use any words or values desired for the ratings. For example, instead of just using low, medium, and high, you can add in low-medium with a value of 2.5 and medium-high with a value of 7.5.

	Low Impact (1)	Medium Impact (5)	High Impact (10)
Low Likelihood (1)	$1 \times 1 = 1$	$1 \times 5 = 5$	$1 \times 10 = 10$
Medium Likelihood (5)	$5 \times 1 = 5$	$5 \times 5 = 25$	$5 \times 10 = 50$
High Likelihood (10)	$10 \times 1 = 10$	$10 \times 5 = 50$	$10 \times 10 = 100$

Table 7-2 Risk Matrix Using Qualitative Analysis

Threat Event	Likelihood	Impact	Risk Score
DoS attack on the intranet web server	Low (1)	Medium (5)	Low-medium (5)
Malware infection	Medium (5)	High (10)	Medium-high (50)
DoS attack on the Internet web server	High (10)	High (10)	High (100)

Table 7-3 Risk Matrix Ratings Used for Adversarial Threat Events

Table 7-3 shows how these ratings can be used for specific threat events. It shows three adversarial threat events with risk scores of 5, 50, and 100. The threat event with a score of 100 (DoS attack on the Internet web server) represents the highest risk and needs to be addressed first.



NOTE Chapter 5 covers various types of attacks, including denial of service (DoS) attacks. A DoS attack attempts to reduce a system's ability to provide a service. In Table 7-2, DoS attacks on the web servers would hamper the web servers' ability to serve web pages.

An advantage of the qualitative analysis is that it provides a clear priority of the different risks. Simply by looking at the result of the analysis, it becomes a simple matter to determine what risks should be addressed first. A disadvantage is that it is difficult to use the results in a cost/benefit analysis when evaluating controls.



EXAM TIP A quantitative analysis uses numerical data such as monetary figures and is objective. A qualitative analysis uses scenario-based data and is subjective. Many organizations use both methods for different risk assessments.

Risk Assessment Steps

Risk assessments can be complex projects taking a lot of time and effort. However, following an outline of predefined steps can simplify them. This also helps ensure that the risk assessment is reproducible and repeatable, and that different security experts can create the same results from the same data. This also helps an organization identify trends between assessments.

With this in mind, it's important to look for previous risk assessments before starting a risk assessment. If they exist on the same system you're evaluating, you can use them as a template; if you're redoing a risk assessment, you only need to identify the changes since the last risk assessment.

An organization isn't required to follow any specific steps. However, NIST SP 800-30 provides four steps that any organization can follow to create reproducible and repeatable results. Figure 7-6 shows the overall steps associated with a risk assessment as recommended in NIST SP 800-30, along with tasks within these steps. The following sections provide additional details on these steps.

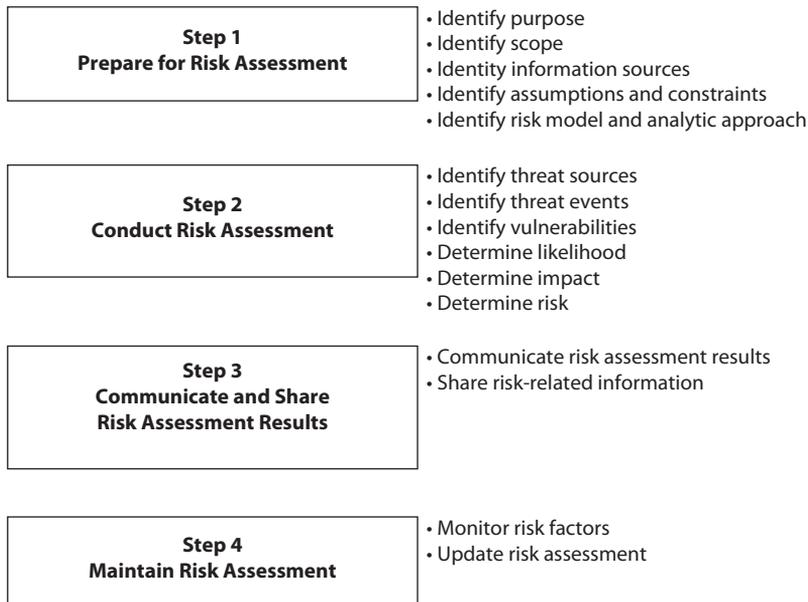


Figure 7-6 Risk assessment steps including tasks within each step

Step 1: Prepare for Risk Assessment

In the first step, the organization performs key tasks to prepare for the risk assessment by identifying several elements used during the assessment. This includes identifying the purpose and scope of the assessment and identifying sources of information for the assessment. Any large project is susceptible to scope creep, where the project slowly grows well beyond the original goals. As an example, if the goal of the assessment is to assess risks associated with e-mail, the study can easily extend into the network infrastructure, including routers and firewalls. It could include power and air conditioning supplied to the entire server room. It could also include malware protection on individual clients and much more. If the goal is to examine only the direct risks associated with the e-mail servers, the purpose and scope should identify this.

This step also includes identifying assumptions and constraints. An organization will typically have safeguards or controls in place designed to mitigate risks. Additionally, some controls may be in the purchasing pipeline but not implemented yet. A risk assessment will often include a list of both current and planned controls, along with clarifying that the risk assessment assumes all of these controls are in place.

Last, this step includes identifying the risk model and analytic approach for the risk assessment. As an example, some risk models refer only to threats, while other models refer to both threat sources and threat events. The two types of analysis are quantitative and qualitative, both of which were covered in the previous two sections.

Step 2: Conduct Risk Assessment

During the risk assessment, you identify threat sources, threat events, and vulnerabilities. You then perform an analysis to determine the likelihood of a risk occurring, the impact if it does occur, and the overall risk.

Threats can be internal or external, natural or manmade, and intentional or accidental. Some models suggest matching threats with vulnerabilities, but this can often be tedious and doesn't always provide any additional actionable data. Many current models (including the one outlined in NIST SP 800-30) do not require matching threat/vulnerability pairs.

After identifying likely threats and vulnerabilities, the assessment then attempts to identify the likelihood of a threat and the impact if it occurs. Security professionals rarely conduct risk assessment themselves. Instead, they often use surveys, questionnaires, and interviews with subject matter experts to gather information. For example, a survey may ask respondents to rank the likelihood of a threat event occurring and the impact if it does occur.



NOTE The specific numbers or words used to identify the likelihood and impact of a risk may be different from one organization to another. A specific scale isn't available that all organizations must use. However, as long as the assessment uses the scale consistently, assessors can choose any scale that works for them.

After identifying the likelihood and impact, it's possible to assign a risk score by combining the likelihood with the impact. A low likelihood combined with a low impact results in a low risk score, while a high likelihood combined with a high impact results in a high score.

The assessment can also assign numbers to the word values. For example, low, medium, and high could be 1, 5, and 10, respectively. It's then possible to calculate the risk score by multiplying the likelihood and the impact. For example, if the probability is 7 and the impact is 8, it provides a risk score of 56.

It's also possible to identify the risk score subjectively. For example, the possibility of a fire is low, while the impact would be high. Multiplying 1 (low) times 10 (high) gives a risk score of 10. However, the assessor can override this relatively low risk score because the impact could be catastrophic.

While a risk assessment typically does not recommend controls, it often evaluates prospective controls or safeguards to determine their value. As mentioned previously, a quantitative analysis makes it easier to evaluate the actual value of a control. However, you can also use a qualitative analysis to evaluate the control. The key is to calculate the risk score without the new control and with the new control.

Step 3: Communicate and Share Risk Assessment Results

After conducting the risk assessment, assessors provide the results to decision makers and other personnel within the organization. Risk assessment results include sensitive information, so the results aren't shared with everyone within the organization. However, if the results directly affect a department, it's common to share the results with department stakeholders.

While this step is typically done after completing the risk assessment, in some cases assessors will provide the results earlier. For example, if the assessment discovers a serious risk, assessors might report this to senior management right away.

The method of communicating and sharing the results will vary from one organization to another. One organization might use informal briefings before finalizing the risk assessment report. Another organization might decide to skip the briefings until the risk assessment report is finalized.

The contents of a risk assessment might also vary from one organization to another. However, you'll often see the following items:

- **Executive summary** This is typically about 5 to 10 percent the size of the full report and summarizes the risk assessment.
- **Risk assessment preparation** This section identifies the purpose, scope, and assumptions identified in step 1 of the assessment. It also documents the information sources, the risk model, and analysis methods used for the assessment.
- **Threat sources, threat events, and vulnerabilities** Depending on the needs of the organization, the assessment might focus on only adversarial threats or all possible threats.
- **Risk scores** The assessment includes the likelihood of the risks, the potential impact, and the risk scores. It's appropriate to document this information in tables or graphs so that it is easy to digest.
- **Controls or countermeasures** If the assessment evaluated controls or countermeasures, it should include details on them. Whenever possible, this section should document the cost of the controls and the predicted impact before and after implementing the controls.
- **Uncertainties** A risk assessment attempts to predict future risks, so it isn't always based on complete data. Instead, it relies on expert opinions, assumptions, and speculations. An uncertainty analysis documents uncertainties caused by missing data, subjective opinions, and assumptions. This helps decision makers evaluate the usefulness of the risk assessment as events change. As an example, advanced persistent threats (APTs) from foreign countries have launched many attacks on companies within the United States. Even though there is some uncertainty about who sponsors specific APTs and their intentions, it is possible to predict future risks associated with APTs.
- **Appendixes** When appropriate, the assessment includes appendixes to document additional information such as details to support any of the conclusions.

Step 4: Maintain Risk Assessment

A risk assessment is a point-in-time report. In other words, it evaluates threats and vulnerabilities at a specific time. However, threats are constantly changing, so an organization must attempt to keep up with the changes with effective risk-management and risk-monitoring decisions.

Organizations periodically repeat or update risk assessments for key systems. The frequency is variable, but it is common to repeat a risk assessment on an annual or semiannual basis.

Maintaining a risk assessment also includes monitoring risk factors identified in the risk assessment. If risk factors change, the organization might choose to update the risk assessment outside of the normal schedule. It certainly isn't necessary to wait for an incident before reassessing risk.

Address Findings

An important step after completing a risk assessment is to address the findings. In many cases, the assessment will identify risks that need to be mitigated by implementing controls. Of course, management makes the decision about which risks to address. In general, a risk assessment does not recommend controls. However, it can evaluate recommended controls and assess their value.

One process that some organizations implement is to periodically review the findings of past assessments and audit the progress. Auditing the findings later (such as six months after the completion of an assessment) provides insight into the accuracy of the predictions. For example, an assessment might have predicted that implementing a specific security control could eliminate a risk. Six months later, you can determine if this was true or not. This step also helps ensure that the organization is implementing the controls approved by management.

Responding to Incidents

In the best of all worlds, all of your risk management practices will prevent any incidents. However, avoiding all incidents is highly unlikely, so organizations also create plans to respond to incidents.

In the context of IT security, a security incident is any violation of policies or security practices that has the potential to result in an adverse event. NIST SP 800-61, *Computer Security Incident Handling Guide*, provides several definitions that are helpful in identifying incidents:

- *An event is any observable occurrence in a system or network.*
- *Adverse events are events with a negative consequence, such as system crashes, network packet floods, unauthorized use of system privileges, unauthorized access to sensitive data, and execution of malicious code that destroys data.*
- *A computer security incident is a violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices.*

Notice that not all events are security incidents. Administrators or security personnel investigate events to determine whether they are incidents. Additionally, not all computer security incidents necessarily result in adverse events. For example, imagine that the company security policy states that users should not use computers to do personal

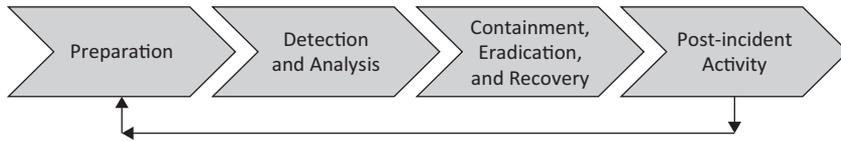


Figure 7-7 Incident response life cycle recommended by NIST SP 800-61

shopping on the Internet. If a user violates this policy, it is an incident, but it won't result in negative consequences, (except perhaps for the employee).

SP 800-61 also provides a standardized process for preparing and responding to incidents. Figure 7-7 shows the life cycle of incident response, and the following sections explore these steps in greater depth.



EXAM TIP The first step in incident response is preparation, which includes creating an incident response plan. After detecting and verifying an incident, it's important to contain the incident as quickly as possible.

The “Identifying Threat Sources” and “Identifying Threat Events” sections earlier in this chapter described many different types of incidents. Additionally, the following list describes some common types of incidents:

- **Denial of service (DoS)** Any type of attack that attempts to reduce a system's ability to provide a service is an incident. This includes both DoS and distributed DoS (DDoS) attacks. Intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) attempt to detect and block DoS and DDoS attacks before they can cause any damage.
- **Malware** The discovery of any type of malicious software, such as a virus, worm, Trojan horse, and so on, is an incident. If AV software discovers the malware before it's installed on a system, it's a minor incident. However, if the malware is able to spread unabated through a network, it can quickly become a major incident.
- **Inappropriate usage** Most organizations have an acceptable use policy identifying what employees can do with IT systems. Any attempt to violate this policy is an incident.
- **Unauthorized access** If any individual or software program is able to access data or systems that the user or program is not authorized to access, it is an incident. Depending on the value of the data or system accessed, the incident may be a minor incident or a major incident.

Preparation

Preparing for an incident is an extremely important first step. With solid preparation, you have a much better chance of minimizing damage from an incident when it occurs. Preparation includes the steps to prevent an incident by ensuring that your systems are secure, but it also includes planning for incidents and planning the response.

Many organizations identify individuals in a computer incident response team (CIRT). These individuals are responsible for responding to an incident when it occurs. A CIRT typically includes members from different parts of the organization and with different skills, such as senior management, information technology, information security, legal, and personnel with public communications skills.



TIP Some organizations use different terms for their teams. CIRT is common, but you may see other terms such as computer security incident response team (CSIRT) or just incident response team (IRT). Each of the terms refers to the same group of people who respond to an incident.

The CIRT needs ready access to tools and resources when an incident occurs. Every minute they have to spend looking for these tools is a minute lost containing the incident, so it's common to locate the tools in an easily accessible location known by all the team members. The following list identifies some of these tools and resources:

- **Contact information** This includes contact information for all CIRT members and key personnel in the organization who may need to be notified of an ongoing incident.
- **Reporting forms** The organization may want to ensure that specific information is recorded about an incident. Pre-created forms help the team members remember what information to document.
- **Forensic tools** Computer forensic tools help personnel collect and analyze evidence while observing relevant requirements to maintain the quality of the evidence. In case the incident results in legal proceedings, it's important for personnel to ensure they don't modify the data as they collect it and that they maintain proper custody of evidence after collecting it. Chapter 13 covers forensic issues in more depth.
- **Documentation** This includes documentation on systems and the network infrastructure. It should include a ready reference of approved changes and the status of these changes. When team members identify something that looks suspicious, such as a system configuration, they need to be able to identify whether it is a valid change or a malicious change. For example, administrators typically disable Trivial File Transfer Protocol (TFTP) when hardening systems, unless it's needed. Some malware enables TFTP to forward files to other systems. If team members see TFTP enabled on a system, this documentation should allow them to verify if it is needed.
- **Software and hardware** The team may need access to hardware to perform simple tasks such as creating reports, doing research, or performing analysis. They also may need access to software used to rebuild systems.
- **War room** In the case of a serious incident, the team may need to use a central war room for coordination. This will likely be something like a general-purpose conference room used only as a war room during an incident. The incident response plan might designate a room to be used for this purpose if necessary.

This is by no means a complete list, and each organization's list will likely be different. However, the important thing to remember is that the CIRT will need ready access to tools and resources. Many of these items could be stored in a network location. However, an incident might take the network down. As an alternative, some organizations create a "crash cart" for incidents. It would include a dedicated computer with relevant forensic tools installed, additional hardware and software that may be needed, and copies of relevant documentation.

Detection and Analysis

The next stage of the incident response life cycle is detection and analysis. At this stage, the incident is discovered, investigated, and analyzed.

There are multiple methods of detecting an incident. It could be as simple as AV software detecting malware on a USB device and sending an alert to notify a user. Users could notice suspicious activity and report it. Systems may randomly lock up or crash. Administrators may see logs growing at an alarming rate, strange files appearing on servers, or simply an increase in network activity (and a slowdown in network response). After detecting these events, someone investigates them to determine whether an incident has occurred.

Intrusion detection systems (IDSs) provide automated detection of potential incidents and display alerts to IT personnel about them. It's important to realize that even though an IDS alerts personnel about activity, that activity isn't necessarily an incident. Consider a SYN flood attack. In day-to-day operations, network issues may prevent the third packet from reaching the server. If this happens once in an hour, it's highly unlikely that it's an attack. Then again, if a server is receiving 100 half-open connections a second, that is very likely an attack.



NOTE Chapter 5 covered SYN flood attacks. As a reminder, they disrupt the TCP handshake process by withholding the third packet to complete the connection. The attacker sends the first packet (the SYN packet), the attacked system responds with a SYN/ACK packet, but the attacker never sends the final ACK packet. This consumes resources on the server and can prevent legitimate users from connecting to the server.

Security professionals who manage an IDS have to define the number of half-open connections that constitutes an attack. Clearly, if this happens once an hour, it isn't an attack. However, if it is happening 100 times a second, it is very likely an attack. They identify the number and set it as a threshold to indicate an event. If they set the threshold too high, they won't be notified when an attack is under way. If they set the threshold too low, they will receive false positives—notifications of possible attacks that are actually not attacks. Given the choice between not being notified of an actual attack or occasionally getting some false positives, most security professionals prefer the latter. They accept some false positives with the goal of ensuring the IDS does notify them of actual attacks.

With this in mind, personnel must investigate and analyze each potential incident. This analysis determines whether the event is an incident and, if so, attempts to prioritize the incident. Minor incidents don't affect critical systems or critical infrastructure. Critical incidents affect mission-critical systems and can seriously degrade the organization's ability to perform its primary mission. Obviously, a critical incident would mandate the use of all available resources to address the incident.

Containment, Eradication, and Recovery

The next stage of the incident life cycle is containment, eradication, and recovery. Although these are included in a single step in SP 800-61, you can actually think of them as three separate steps.

Once personnel verify that an event is an incident, the next step is to contain the incident as quickly as possible. The goal is to prevent the incident from spreading. For example, after malware infects a single system on a network, it typically tries to locate and infect other systems on the network. However, disconnecting the network cable on the infected system contains the malware to just that system. It cannot spread on the network. In contrast, if you didn't contain the malware, it could quickly spread to all the computers on the network, causing a major incident.



EXAM TIP Once an incident has been identified, it should be contained as soon as possible. This can't occur before detecting the incident, but you can contain an incident after you've verified that it is an actual incident and not a false positive.

Once you've contained the incident, you move to eradication. For example, some multipartite viruses have multiple components, and you must remove each individual component. If you miss any component, the malware can reemerge and reinfect all the files you've cleaned. Some malware can reconfigure systems, weaken security by enabling services or protocols, create new accounts used for backdoor access, or open previously closed ports on firewalls. All of the actions taken by the malware needs to be reversed to eradicate the malware.

The last step in this stage is recovery, and what it entails depends on the damage caused by the incident. For minor incidents that haven't caused any damage, recovery may not require any steps at all. However, a major incident may require the recovery of an entire server or even an entire location. For example, if a fire destroyed a building, recovery entails activating an alternative location, moving critical systems and data to the alternative location, and bringing everything back online.



NOTE Chapter 12 covers disaster recovery operations within the context of disaster recovery plans. It also covers alternative locations such as hot sites, cold sites, warm sites, and mobile sites. Each of these sites provides varying levels of readiness for disaster recovery.

Post-incident Activity

The last step is post-incident activity. In this step, you examine the incident and the response. The goal is to determine whether the existing plans and procedures adequately addressed the incident or if there are response procedures that need improvement. Organizations that take the time to examine their response typically find ways to improve their existing processes.

One of the ways to perform a post-incident activity is by holding a meeting with people involved in the incident and posing questions to generate discussions. Obviously, you wouldn't hold a meeting like this for a minor incident, but if an incident negatively affected the mission of the organization, a post-incident activity meeting is warranted. The information gathered during this meeting is compiled into a post-incident report that can be used in future risk management activities such as risk assessments.

The following are some of the questions you can include in this meeting:

- Are there any lessons to be learned from this incident? What are they?
- Did we respond to the incident as quickly as possible? If not, what can we do to improve the response time for a subsequent incident?
- Was there anything that prevented the recovery from occurring in a timely manner?
- Was the documentation up to date? If not, what was not up to date? Is there a way that we can ensure it is up to date in the future?
- Did the CIRT have all the tools it needed? If not, what tools were not available? How can we ensure these tools are available if we have a similar incident?
- Are there any preventive measures we can implement either to prevent a similar incident or to at least reduce the severity of the incident?

One of the challenges of performing a post-incident activity meeting is to avoid finger pointing and “blame storming.” Instead of focusing on improving the process for a potential event in the future, some participants may focus on ensuring that they aren't blamed for this event. If senior management stresses that the goal is to improve the process and not to place blame, participants may be more willing to provide constructive input. Of course, if management uses the post-incident report to hold individuals or departments accountable, it's highly unlikely future post-incident activity meetings will have any success.

Chapter Review

Risk is the probability or likelihood that a threat will exploit a vulnerability and cause a loss. It is sometimes represented in the formula of $\text{Risk} = \text{Threat} \times \text{Vulnerability}$. A threat is any activity that can be a possible danger and threat events include any type of activity or event that can result in a loss of confidentiality, integrity, or availability of an organization's assets. A vulnerability is a weakness, and a loss represents a negative event for an organization.

NIST publishes several documents related to IT security, and two important documents related to risk are SP 800-30 (*Guide for Conducting Risk Assessments*) and SP 800-61 (*Computer Security Incident Handling Guide*).

Risk management includes several elements used to reduce risk to a manageable level. This includes methods to mitigate (or reduce) risks, avoid risks, share or transfer risk, and accept risk. Residual risk is the risk that remains after mitigating risks. Senior management is responsible for any losses resulting from risk they choose to accept, including residual risk.

One of the first steps in risk management is to identify assets. Once an organization identifies the assets that it considers valuable, it then takes steps to protect them. An organization addresses high risks by implementing security controls. Security controls typically attempt to reduce vulnerabilities or reduce the impact of a risk. However, if the valuable assets haven't been identified, it's possible that they won't be protected and that resources will be expended to protect less valuable assets.

A risk assessment is a point-in-time evaluation. It examines a system to identify threats and vulnerabilities, based on the current security controls. It then attempts to evaluate the likelihood of a threat exploiting a vulnerability and the impact, or magnitude of harm, from this event. Identifying the likelihood and the impact helps to determine the level of risk. Risk management processes address high risks.

Risk assessments commonly use either a quantitative analysis or a qualitative analysis. A quantitative analysis uses numerical figures to calculate risk. The single loss expectancy (SLE) identifies the loss from a single event. Annual rate of occurrence (ARO) indicates how many times you expect the event to occur. Annual loss expectancy (ALE) provides the total expected loss for a year and is calculated as $SLE \times ARO$. This is useful to evaluate the value of a control. If the cost of the control is less than the ALE, it is justified. If it is significantly higher than the ALE, it is not justified. If the cost of the control and the ALE are relatively close, they must be evaluated to determine a long-term return on investment (ROI).

Many organizations designate a computer incident response team (CIRT or IRT) to respond to incidents. It's important to know the different steps in incident response. SP 800-61 defines these steps as (1) preparation; (2) detection and analysis; (3) containment, eradication, and recovery; and (4) post-incident activity.

Questions

1. Which of the following choices best represents the definition of risk?
 - A. The likelihood that a threat source can cause a threat event resulting in a vulnerability
 - B. The likelihood that a vulnerability can exploit a threat and cause a loss
 - C. The likelihood that a threat will exploit a vulnerability and cause a loss
 - D. The likelihood that an incident can cause a vulnerability resulting in a loss

2. You are involved in risk management activities within your organization. Of the following activities, which one is the best choice to reduce risk?
 - A. Reducing threats
 - B. Increasing vulnerabilities
 - C. Increasing impact
 - D. Mitigating risk
3. Which of the following choices identify valid threat sources? (Choose all that apply.)
 - A. Employee
 - B. Earthquake
 - C. State-sponsored attacker
 - D. Administrator
4. Which of the following can cause a negative impact on an organization's assets?
 - A. A threat
 - B. A risk
 - C. A weakness
 - D. A control
5. What's a primary method used to reduce risk?
 - A. Reducing threats
 - B. Reducing vulnerabilities
 - C. Increasing threats
 - D. Increasing vulnerabilities
6. What is the purpose of risk management?
 - A. Eliminate risks
 - B. Reduce risks to an acceptable level
 - C. Share or transfer risks
 - D. Identify risks
7. What should an organization do when the cost of a control exceeds the cost of a risk?
 - A. Implement the control
 - B. Accept the risk
 - C. Perform a risk assessment
 - D. Mitigate the risk

8. Which of the following helps ensure that an organization focuses risk management resources only on the most serious risks?
 - A. Risk assessment
 - B. Residual risk
 - C. Countermeasures
 - D. Qualitative analysis
9. You decide to manage risk by purchasing insurance to cover any losses. Which one of the following risk management techniques are you using?
 - A. Accept
 - B. Avoid
 - C. Mitigate
 - D. Transfer
10. An organization has a business location in Miami, Florida. Due to the risks associated with hurricanes, the organization has decided to move the location to Atlanta, Georgia, away from any ocean. What risk management strategy is the organization using?
 - A. Accept
 - B. Avoid
 - C. Mitigate
 - D. Transfer
11. An organization has implemented several controls to mitigate risks. However, some risk remains. What is the name of the remaining risk?
 - A. Vulnerable risk
 - B. Mitigated risk
 - C. Alternate risk
 - D. Residual risk
12. A risk assessment recommended several controls to mitigate risks, but only some of the controls were accepted and implemented. Who is responsible for any losses that occur from the remaining risk?
 - A. The person completing the risk assessment
 - B. Senior management
 - C. IT personnel managing the systems
 - D. Security personnel

13. Of the following choices, which one most accurately reflects differences between risk management and a risk assessment?
 - A. A risk assessment is a point-in-time event, while risk management is an ongoing process.
 - B. Risk management is a point-in-time event, while a risk assessment is an ongoing process.
 - C. Risk assessments are broad in scope, while risk management is focused on a specific system.
 - D. Risk management is one part of an overall risk assessment strategy for an organization.
14. Of the following choices, what is an important first step in a risk management plan?
 - A. Implementing controls
 - B. Identifying vulnerabilities
 - C. Identifying assets
 - D. Identifying threats
15. You are completing a risk assessment using historical data. You've identified that a system has failed three times in the past year, and each of these outages resulted in approximately \$10,000 in losses. What type of analysis does this allow you to perform?
 - A. Qualitative
 - B. Quantitative
 - C. Informative
 - D. Subjective
16. Which of the following formulas will determine the annual loss expectancy (ALE)?
 - A. $SLE - ARO$
 - B. $SLE \times ARO$
 - C. $ARO - SLE$
 - D. SLE / ARO
17. You are completing a risk assessment and using historical data. You've identified that a system has failed five times in each of the past two years, and each outage resulted in losses of about \$5,000. What is the ARO?
 - A. Five
 - B. \$5,000
 - C. \$25,000
 - D. Impossible to determine with the information provided

18. You have completed a risk assessment and determined that you can purchase a control to mitigate a risk for only \$10,000. The SLE is \$2,000 and the ARO is 20. Is this cost justified?
- A. Yes. The control is less than the ALE.
 - B. No. The control exceeds the ALE.
 - C. Yes. The control exceeds the ARO.
 - D. No. The control is less than the ARO.
19. What is the first step in incident response?
- A. Analysis
 - B. Containment, eradication, and recovery
 - C. Detection
 - D. Preparation
20. Of the following choices, what best represents all of the steps related to incident response?
- A. Preparation, containment, detection, analysis, eradication, and recovery
 - B. Preparation, detection, analysis, containment, eradication, and recovery
 - C. Containment, preparation, detection, analysis, eradication, and recovery
 - D. Containment, analysis, detection, eradication, and recovery

Answers

1. **C.** A risk is the likelihood that a threat will exploit a vulnerability and cause a loss. Threats do not create vulnerabilities, and vulnerabilities do not exploit threats. Similarly, incidents do not cause vulnerabilities.
2. **D.** Risk mitigation is the process of reducing risk. You can rarely reduce threats, but you can often reduce (not increase) vulnerabilities or reduce (not increase) the impact of a risk.
3. **A, B, C, D.** All of the answers are valid threat sources. Employees can be adversarial threat sources if they intentionally cause damage or accidental threat sources if they accidentally cause damage. An earthquake is an environmental threat source. A state-sponsored attacker is an adversarial threat source.
4. **A.** A threat source can cause a negative impact by exploiting a vulnerability. Risk is the likelihood that a threat will exploit a vulnerability and cause a loss; the risk doesn't cause the negative impact. A weakness is a vulnerability. Attackers can exploit a vulnerability, but the vulnerability doesn't cause the loss. Controls attempt to reduce risk by reducing vulnerabilities or reducing the impact of a risk.

5. **B.** A primary method of risk mitigation is reducing vulnerabilities. Threats often can't be reduced, and adding more threats won't reduce risk. You reduce vulnerabilities by implementing controls.
6. **B.** Risk management reduces risks to an acceptable level. It is not possible to eliminate risk. One method of managing risk is to share or transfer risk, but that is not the only method. Similarly, risk management processes identify risk, but risk management is much more than just identifying risk.
7. **B.** If the cost of a control exceeds the cost of a risk, the organization should accept the risk. The organization might implement the control if the cost of the control was less than the cost of the risk, indicating a cost savings. A risk assessment can analyze the value of the control, but you wouldn't need to do a risk assessment if you already know the cost of the control exceeds the costs of the risk. Mitigating the risk indicates you would implement the control, but based on the known costs, it's appropriate to accept the risk.
8. **A.** A risk assessment helps an organization prioritize risks so that it can focus risk management resources on the most serious risks. Residual risk is the risk that remains after implementing risk mitigation steps. Countermeasures are risk management resources that reduce risks. Risk assessments can use either a quantitative analysis or a qualitative analysis, but neither is superior to the other one in all circumstances.
9. **D.** Insurance is one of the ways that you can manage risk by transferring the risk to a third party. Risk acceptance doesn't take any further action to mitigate the risk. In risk avoidance, you avoid the activity that results in the risk. It's most common to try to reduce the risk using risk mitigation.
10. **B.** By moving the location to a city that can't be hit by a hurricane, the company is using risk avoidance. Risk acceptance doesn't take any action to mitigate the risk. In risk mitigation, you attempt to reduce the risk, perhaps by ensuring that the building is built with hurricane-resistant materials. The company can transfer the risk by purchasing hurricane and flood insurance.
11. **D.** Residual risk is any risk that remains after implementing controls to mitigate the risk. It's often not cost effective to implement controls to eliminate all risks, so senior management must make decisions on what risk to mitigate and what risk to accept as residual risk. Vulnerable risk and alternate risk are not valid terms associated with risk management.
12. **B.** Senior management is responsible for making decisions on what risk to mitigate. The remaining risk is residual risk, and senior management is responsible for any losses from this residual risk.
13. **A.** A risk assessment is a point-in-time event, while risk management is an ongoing process. Risk assessment is one element of a risk management strategy, and risk assessments are generally focused on specific systems with a limited scope, while risk management is much broader.

14. **C.** You must identify assets first. You can then identify threats against these assets and vulnerabilities in these assets. You can't recommend or implement controls until you know what you want to control.
15. **B.** A quantitative analysis uses numerical figures to identify the actual costs associated with a risk. A qualitative analysis uses subjective terms such as low, medium, and high to analyze a risk. There is no such thing as an informative analysis.
16. **B.** ALE is the product calculated from the single loss expectancy (SLE) and the annual rate of occurrence (ARO), or $SLE \times ARO$. The ALE is not calculated by subtraction of the ARO and SLE or by dividing the SLE and ARO.
17. **A.** The annual rate of occurrence (ARO) is five because it happened five times each in the past two years. The single loss expectancy (SLE) is \$5,000 and the annual loss expectancy (ALE) is \$25,000.
18. **A.** Because the cost of the control is less than the annual loss expectancy (ALE), the cost is justified. The cost of the control is \$10,000 and the ALE is \$40,000. The annual rate of occurrence (ARO) is how many times the loss occurred (20 in the example), but it is only useful when you multiply it with the single loss expectancy (SLE) to identify the ALE.
19. **D.** The first step in incident response is preparation, which includes creating an incident response plan. The other answers are valid steps in incident response, but they aren't the first step.
20. **B.** The steps recommended in NIST SP 800-61 are preparation, detection, analysis, containment, eradication, and recovery. Containment is important once an incident has been detected and analyzed, but can't be done beforehand.

Monitoring and Analysis

In this chapter, you will learn about

- Intrusion detection systems
- Security alerts and false positives
- Differences between network-based and host-based intrusion detection systems
- Differences between intrusion detection systems and intrusion prevention systems
- Methods of detecting and preventing attacks
- File integrity checkers
- Honey pots, honeynets, and padded cells
- Security event management tools used in large networks
- Vulnerability assessment test types
- Vulnerability assessment tools
- Penetration testing

Operating and Maintaining Monitoring Systems

Monitoring systems within an information technology (IT) network helps to prevent, detect, and correct potential security incidents. Intrusion detection systems provide continuous monitoring of a network and report events of interest that might be an incident. Intrusion prevention systems go a step further and prevent malicious traffic from entering a network, thereby preventing potential incidents. The following sections describe these systems in more depth.

Intrusion Detection Systems

Intrusion detection systems (IDSs) provide continuous monitoring of networks and hosts to help protect them from attacks. The goal is to detect an attack as it occurs and raise an alert so that administrators can take steps to stop the attack. Some IDSs are passive and some are active:

- **Passive** The IDS monitors and records events in a log. It provides alerts (such as via e-mail or a text message) when it detects a potential attack.

- **Active (or reactive)** The IDS takes action to modify the environment to stop the attack and reduce its effects. For example, it might modify an access control list (ACL) on a router to block traffic or close connections on a server. An active system also logs events and sends alerts to notify personnel of potentially malicious events.

IDSs are either network-based or host-based. A network-based IDS (NIDS) monitors network traffic, while a host-based IDS (HIDS) monitors traffic for an individual system such as a server or workstation.

Figure 8-1 shows a network with both a NIDS and a HIDS. Notice that NIDS agents are located on different network segments, such as before the DMZ, inside the DMZ, and on the internal network. HIDS software is installed on each individual server and workstation.

A NIDS collects information from the NIDS agents. Administrators install these agents on servers, routers, firewalls, or any other device that can monitor potentially malicious traffic coming into the network. In some cases, the NIDS agent forwards logs from these event source systems, and in other cases, the NIDS agent creates its own log and forwards it to the NIDS management system.

IDS Alerts

When an IDS detects an attack, it logs it and typically reports it as an alert or an alarm. These reports come in many forms. Many IDSs include the ability to send an alert via an e-mail or a text message.

It's also possible to display the alert. As an example, a network operations center (NOC) typically has a wall of monitors in front of administrators and technicians. The administrators and technicians work on various tasks with desktop computers, but frequently glance at the monitors, so alerts displayed there will get their attention.

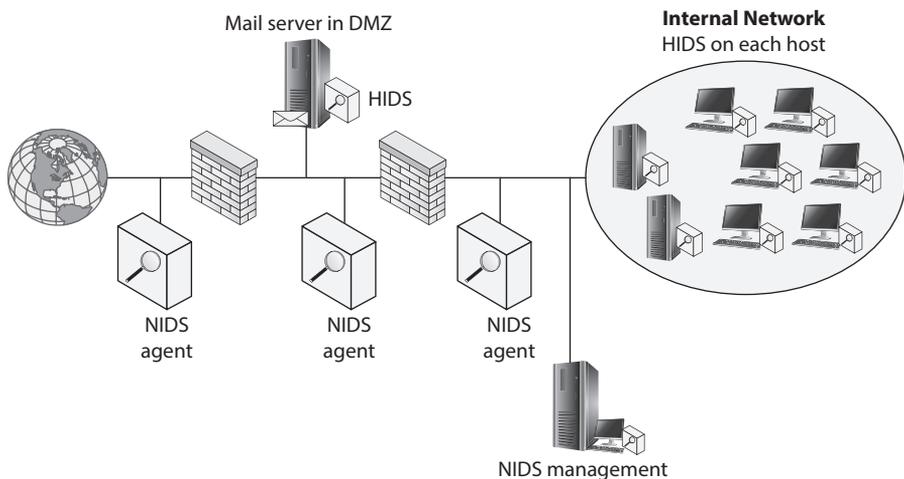


Figure 8-1 Intrusion detection systems protecting a network and systems

An alert or an alarm doesn't necessarily indicate an attack is under way. Personnel investigate the alerts to determine if they are valid. Unfortunately, it's common for an IDS to give false positives. In other words, an IDS might indicate a possible attack when an attack isn't underway at all.



EXAM TIP An alert provides notification of a potential adverse event. Personnel analyze the event to determine if it is an incident. An alert can be a false positive, which isn't an incident.

An IDS often triggers an alert when an event reaches a specific threshold. As an example, consider a port scan attack, where an attacker attempts to scan a system's ports to identify open ports. If the scan detects that port 80 is open, the attacker knows that the system is probably a web server running HyperText Transport Protocol (HTTP) because port 80 is the well-known port for HTTP. A port scan will scan a list of ports and record what ports elicited a response (and were open) and what ports did not elicit a response.

If a remote system scans one port, that is probably not an attack. However, if an unknown external system scans all 1,024 well-known ports in a 60-minute period, it is very likely an attack.

Here's a trick question. If one port scan in an hour is not an attack, and 1,024 port scans in an hour is an attack, what is the lowest number between 1 and 1,024 that most likely indicates an attack? In other words, what should you set the threshold to so that the IDS detects a port scan attack?

There just isn't a good answer to that question. Port scanners allow attackers to randomize the ports they scan and set delays between queries. If the attacker sets the delay to five minutes, the scanner does 12 port scans in an hour. If you set the threshold to 15, your system would be under attack but the IDS would not detect it.

In contrast, if you configure the threshold to 2 so that the IDS sends an alert if it detects two port scans in a hour, it will probably create many false positives. Your IDS will be known as the IDS that cries wolf and administrators will ignore it. Many administrators consider a threshold of two port scans in a minute to be too low because it will create so many false positives.

When choosing a number for this type of threshold, administrators recognize that the IDS will probably generate some false positives. They would rather see some false positives than configure the threshold so high that the IDS doesn't detect attacks in progress.

Network-based Intrusion Detection Systems

A NIDS monitors the network traffic for any type of attack. It typically has several nodes or agents stationed around the network installed on routers, firewalls, and network appliances such as unified threat management (UTM) devices. Each of these nodes monitors the traffic and reports its findings to a NIDS management server.

Figure 8-2 shows one possible configuration. It has NIDS agents installed on two firewalls and one router, and each of these agents reports events to the NIDS management server. The server monitors the traffic and, if it detects an attack, provides an alert to administrator or security personnel.

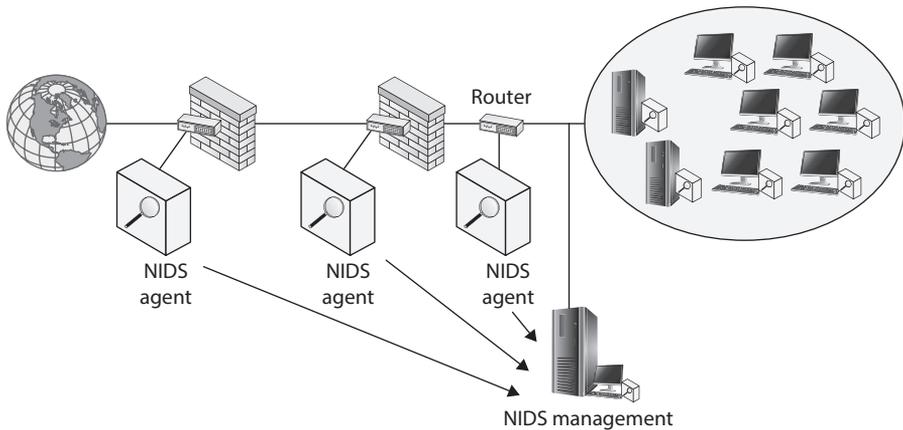


Figure 8-2 Network-based intrusion detection system with a NIDS management server

As discussed in Chapter 4, many routers and switches support port mirroring. As a reminder, port mirroring sends a copy of all traffic passing through a device to a specific port, and this is one way for agents to send traffic to a NIDS management server. Administrators configure the device to use port mirroring and connect the NIDS management server to the mirrored port.

One of the primary benefits of a NIDS is that the monitoring is transparent. Attackers don't have any clear indication that a NIDS is monitoring their activity. In contrast, it's easier for attackers to detect and disable a HIDS.



EXAM TIP A network-based IDS (NIDS) monitors traffic going through a network. It uses agents to monitor traffic on network devices, and the agents forward the traffic to a central management console.

Host-based Intrusion Detection Systems

A HIDS is installed on an individual system such as a server or workstation. It primarily monitors activity on the system such as system processes and running applications. It can only monitor activity on the host; it is not able to monitor overall network activity as a NIDS does. However, a HIDS can look at patterns on a host's disk and within its random access memory (RAM) in ways that a NIDS cannot.

Some organizations install a HIDS on every system in the network, while other organizations install them only on specific systems such as critical servers. The key is that a HIDS monitors systems only, while NIDS agents are installed on network components and monitor network activity.

There are some drawbacks associated with a HIDS. First, the software can be very processor-intensive, slowing down an otherwise powerful computer. Second, installing a HIDS on each of your internal systems can be quite expensive. Last, many attackers can detect a HIDS, then disable it or erase the HIDS's logs.

IDS and Defense-in-Depth

IDSs provide another layer of protection to systems and networks and are part of an overall defense-in-depth strategy. Firewalls monitor traffic going in and out of a network and can block much of the malicious traffic. Antivirus (AV) software detects malware and cleans, deletes, or quarantines malicious files. The IDS focuses on intrusion attacks and can provide protection that is beyond the capabilities of the firewall and AV software.

Chapters 4 and 6 discuss unified threat management (UTM) appliances and content-filtering appliances, respectively, and many of these appliances include intrusion detection components. They can also include spam filtering, malware protection, firewall services, proxy server components, and intrusion detection protection. A unified solution in a single appliance is typically easier to administer and manage.

Additionally, using HIDSs with a NIDS is part of a defense-in-depth strategy. This is similar to installing AV software on each system even though AV software on a UTM appliance monitors all traffic coming into the network. The extra AV protection protects against a threat coming from another source, such as a USB drive inserted into a system. Similarly, installing a HIDS on each system protects against an attack launched internally, either from an internal attacker or from malware inadvertently released on an internal network.



EXAM TIP A host-based IDS is installed on a system such as a server or workstation. It monitors activity on the host, but cannot monitor network activity.

Intrusion Prevention Systems

An intrusion prevention system (IPS) is an advanced implementation of an IDS. However, IPSs have been around long enough that they are often lumped together with IDSs. For example, National Institute of Standards and Technology (NIST) Special Publication (SP) 800-94 is titled *Guide to Intrusion Detection and Prevention Systems (IDPS)*.

However, there is a key difference between an IDS and an IPS. Specifically, a network-based IPS is placed inline with the traffic. For example, Figure 8-3 shows an IPS placed at the boundary between the network and the Internet. It monitors all traffic going in and out of a network and actively blocks potentially malicious traffic. All IPSs are active, while IDSs can be either passive or active.



EXAM TIP A network-based IPS is placed inline with the traffic and includes both passive and active components. It has the capability of modifying the environment to block an attack while it is in progress. An active IDS is often referred to as an IPS. However, a true IPS is placed inline with the traffic and prevents the traffic from entering the network. This isn't true with all active IDSs.

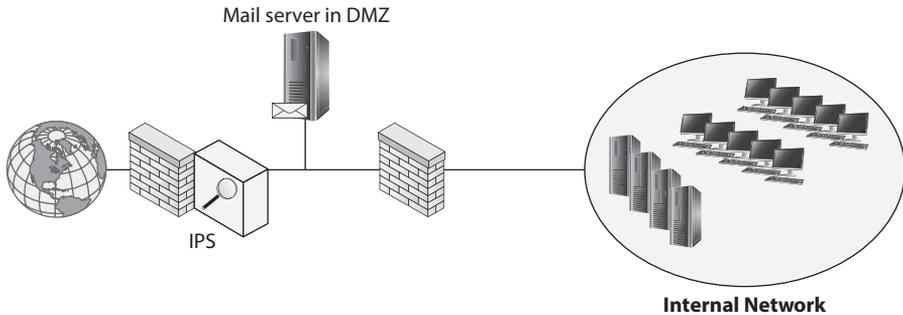


Figure 8-3 Intrusion prevention system placed inline

Although Figure 8-3 doesn't show additional IDSs such as HIDSs on individual systems, they can be used with an IPS as part of a defense-in-depth strategy. An IPS can be host-based (as a HIPS) or network-based (as a NIPS), just as a traditional IDS can be host-based or network-based.

Notice that the IPS is placed inline with the firewall and that all traffic to and from the network goes through both the firewall and the IPS. The firewall protects a network from some malicious traffic, and the IPS helps protect the network from attacks that get through the firewall. Just as an active IDS can modify the environment, an IPS can also modify the environment. For example, it might modify the ACL on a router if it detects a port scan from a specific IP address.



NOTE Many IDSs have the capability of providing an active response and modifying the environment. Such systems are sometimes marketed as an IDS and sometimes marketed as an IDS/IPS or IDPS.

As another example, consider Figure 8-4. It shows a server being attacked with a SYN flood attack. Chapter 5 covers SYN flood attacks, but as a reminder, an attacker disrupts the Transmission Control Protocol (TCP) three-way handshake process, leaving half-open connections on a server. The attacker sends the first SYN packet, the server responds with the SYN/ACK packet, and then the server waits for the ACK packet. However, the attacker never replies with the ACK packet. The server can have hundreds of half-open connections waiting for the third packet in each.

The IPS can detect this attack and block any more traffic from the attacker by modifying the ACL on the router to stop the attack. It can also close the half-open connections to free up the resources on the server that are being consumed by the attack, as shown in Figure 8-5.



NOTE Modifying the ACL on a router doesn't necessarily stop a SYN flood attack. Attackers often use botnets to launch distributed denial of service (DDoS) attacks. A protracted attack could be coming from tens of thousands of computers, none of which is the actual attacker.

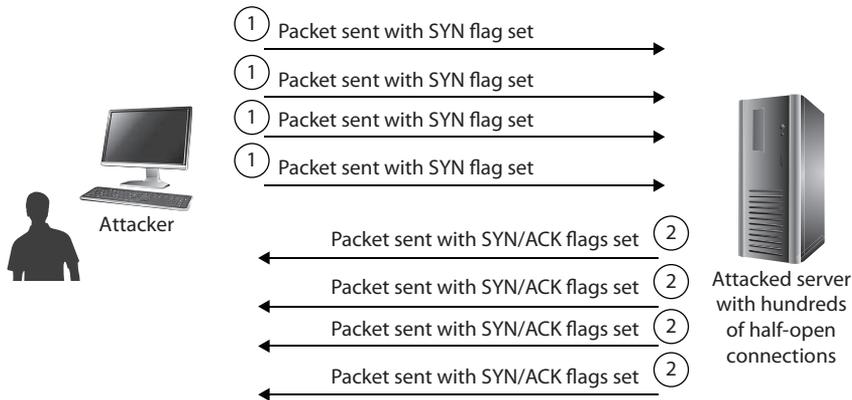


Figure 8-4 SYN flood attack

It's worth emphasizing the active component of an IPS and comparing it to the passive component of an IDS. A passive IDS logs the event and provides notification. All IDSs and IPSs have this ability. The IPS (or active IDS) takes action to modify the environment to stop the attack.

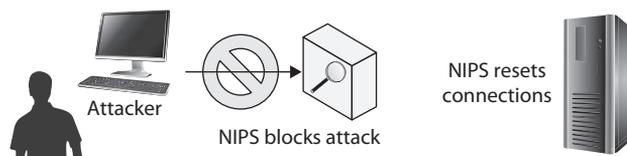
Also, it's important to realize that neither an IDS nor an IPS launches attacks against an attacker. An IPS will change the environment to block the attack, but does not take any aggressive action against the attacker. Many organizations have security professionals with specialized skills who follow attackers through the Internet and launch attacks on them, but these skills are well beyond the scope of basic IT security professionals.

One of the basic risks of attacking an attacker is drawing attention to your organization. If an attacker attempts an attack and fails, that attacker may simply move on to another, easier target. However, if an attacker is attacked by an organization, that attacker may take it personally and launch protracted attacks against the company, holding a grudge for years. Additionally, there may be legal or liability issues with an organization launching an attack on a suspected attacker. Worse, because many attackers spoof their IP addresses or launch attacks through zombies controlled via a botnet, you may be attacking an unsuspecting victim.



EXAM TIP An intrusion detection system (IDS) is a detective security control. An intrusion prevention system (IPS) is a preventive security control.

Figure 8-5
IPS blocking a
SYN flood attack



Detection Methods

IDSs and IPSs use one of the following types of detection methods to detect attacks: signature-based, anomaly-based, or a hybrid of the two. Any type of IDS or IPS can use these methods.



EXAM TIP The two methods of intrusion detection are signature-based and anomaly-based detection. These are sometimes called knowledge-based (using signatures) and behavior-based (using a baseline to detect anomalies).

Signature-based

A signature-based (sometimes called knowledge-based) detection method is similar in concept to AV signatures. Recall that malware has unique characteristics, called signatures, that identify them. AV software has a database of signatures and compares potential malware against the database. If it gets a match, the software identifies the malware and attempts to clean or delete it.

Similarly, many attacks have unique characteristics documented in signature files. The IDS/IPS uses these signature files to identify and detect attacks. Both the vendor supplying the IDS and then the organization using the IDS must update these signature files regularly to detect newer attacks.

This shows a primary drawback with signature-based detection: It's good with known attack tools, but attacks may go unnoticed if an attacker uses modified attack methods.

Anomaly-based

Another method of detection is anomaly-based. Administrators configure the device to monitor normal behavior to create a baseline. The device then observes activity and compares it to the baseline. If the current activity differs significantly from the baseline, the device issues an alert on the activity.

Anomaly-based systems go by many names, including behavior-based and heuristics-based. The important point is that they start by identifying normal activity and documenting this in a baseline, and then they detect abnormal activity that is outside the normal range of the baseline.

Although artificial intelligence has advanced quite a bit in recent years, anomaly-based systems haven't quite achieved that level of sophistication. The result is that an anomaly-based system often has a higher level of false positives than a signature-based system. For example, if an organization has a successful e-mail campaign resulting in a large customer response, this increased customer activity could be interpreted as an attack.

If the monitored system is modified, it's important to update the baseline. For example, if you are using a NIDS and you install software that increases network activity, the NIDS might generate many false positives if the baseline isn't updated.



EXAM TIP If the monitored environment is modified, it's critical to update the baseline in an anomaly-based system. Otherwise, the modification will be detected as an anomaly, resulting in a higher level of false positives.

Hybrid

A hybrid detection method simply uses a combination of both signature-based and anomaly-based detection methods. It's common for many IDSs and IPSs to use a combination of both detection methods to increase their overall performance.

Whitelisting

Although it isn't formally identified as a detection method for IDSs, the use of whitelisting is rising as an alternative to host-based IDSs. The whitelist detects abnormal activity by specifically identifying authorized activity. For example, an application whitelist identifies approved applications and processes. A system blocks applications and processes that aren't on the list.

Similarly, a switch can have a listing of system media access control (MAC) addresses authorized to communicate through a switch. Any systems not on the list that attempt to communicate through the path raise an alarm. This strategy isn't feasible for a large number of computers, but it can be useful for smaller networks. Still, it's important to realize that attackers can spoof MAC addresses. Whitelisting can be part of an overall strategy, but shouldn't be a primary method.

Wireless Intrusion Detection and Prevention Systems

A wireless intrusion detection system (WIDS) works similarly to an IDS, but for wireless networks, and a wireless intrusion prevention system (WIPS) works similarly to an IPS, but for wireless networks. They monitor wireless network traffic, looking for suspicious activity.

Chapter 5 mentions WIDS and WIPS as countermeasures against wireless attacks. They are effective at detecting common attacks, such as from rogue access points, evil twins, and MAC spoofing. They use signature-based detection methods to detect traffic patterns from known attacks, and most include anomaly-based detection methods too. Additionally, they typically have the ability to inspect access points and identify configuration problems.

Analyzing Results

An important element of continuous monitoring is regular analysis of the monitor results. Administrators commonly use automated tools to assist in the analysis. The following list describes some common methods of analyzing the data:

- **Analyzing individual events** Individual events must be analyzed to determine if they are actual incidents. Administrators often use a variety of tools in the analysis. This includes logs on servers, routers, and firewalls in addition to logs collected by the IDS or IPS. In some cases, administrators need to collect additional data, such as by using a sniffer to capture and analyze the packets during an active attack, or the memory contents of an attacked system.
- **Reviewing security analytics and metrics** Administrators attempt to pull relevant data from the logs to create actionable information. For example, it's

possible to measure the number of false positives for each week. These findings may show that the alert threshold is too high or too low and needs to be adjusted.

- **Identifying trends** Metrics help identify a baseline and regular changes to that baseline. For example, an organization might be experiencing a steady increase in denial of service attacks. Analysts can compare these attacks to performance issues with the attacked servers to determine if the security controls are providing adequate protection.
- **Creating graphics for visualization** You don't necessarily need pictures, but many tools allow you to create graphs and other graphics that are easily worth a thousand words. It is much easier to identify trends based on collected metrics with a graphic than it is to wade through text.
- **Communicating findings** After completing an analysis, it's important to communicate the findings. Organizations have different processes in place for this. Some organizations might have personnel use informal reports during the analysis, followed by a formal report after the analysis is complete.

Detection Systems and Logs

One of the challenges of detecting attacks is that attackers often try to erase their tracks. Logs will record what the attacker has done, but if the attacker is able to delete or modify the logs, the record is lost. With this in mind, it's important to realize that any logs on a local system that has been attacked are suspect. You need to investigate to determine whether the logs were modified during the attack.

On the other hand, logs stored on remote systems are the most valuable after an attack. For example, NIDS agents send a copy of their logs to a central NIDS management server. If a network firewall or resources within a DMZ are attacked, the logs on those resources may have been modified, but the logs on the NIDS server will still be intact and valid.



EXAM TIP After an attack, remote logs are the most valuable to re-create the events that occurred during and prior to the attack. Any logs on a local system should be treated with suspicion because the attacker may have modified them.

Detecting Unauthorized Changes

One of the ways that an IDS can detect an attack is by detecting unauthorized changes. In some cases, the changes can indicate a loss of integrity, and in other cases, they can indicate an attack in progress.

File Integrity Checkers

Many IDSs and antivirus applications use file integrity checkers to detect unauthorized changes to critical system files. A file integrity checker uses hashing techniques to detect these changes.

Chapter 1 introduces hashing and Chapter 14 covers it in more depth. You may remember that a hash is simply a number. A hashing algorithm creates the hash when it is executed against data, and no matter how many times the algorithm is executed against the same data, it will always create the same hash. However, if the data is modified, a new calculated hash will be different from the original hash. By comparing hashes created at different times, it's possible to determine whether the original file has lost integrity.

A file integrity checker calculates hashes on critical system files to establish a baseline. It then periodically recalculates the hashes of these files and compares the two hashes. If the hashes are different, it detects the change and sends an alert.

Of course, some files are modified as a part of normal system operation. For example, users modify their own documents regularly, so the file integrity checker doesn't monitor these files. However, it's rare for executable files to be modified except through authorized updates. If the file integrity checker identifies changes in executable files, it often indicates a problem such as a malware infection.

Unauthorized Connections

Many IDSs monitor the connections on systems and attempt to detect unauthorized connections. For example, if a server is hosted in an extranet and should be accessed only by a partner organization, an IDS can detect any unauthorized connections from unauthorized entities.

If a server is hosted within a DMZ and accessible to anyone on the Internet, detection of unauthorized connections isn't so easy. However, an IDS can identify connection attempts that are outside the scope of the server's purpose. For example, it's normal for a web server to have connections using HTTP and HTTP Secure (HTTPS). However, connection attempts using other protocols such as Telnet or Simple Mail Transport Protocol (SMTP) may not be authorized and indicate a potential incident.

Attackers often use scanning tools to learn information about a system, and these scanning tools often give telltale signs of their intentions. An IDS can detect these scans and send an alert.

Honeytrap

A honeypot is a computer set up to entice would-be attackers. Administrators copy dummy data onto the server and configure it with weak security so that an attacker can easily hack into it. In essence, it's a trap designed to look appealing to attackers, just as a pot of honey looks appealing to Pooh Bear of *Winnie the Pooh* fame.

As an example, an organization can set up an FTP server with weak security that looks like it is hosting company data. Attackers may see the server, bypass the weak security controls, and attempt to download or upload data. While the attacker is attacking and probing the honeypot, that attacker is not in the actual network. Additionally, while the attacker is attacking, security professionals have an opportunity to observe the tactics used by the attacker. This observation capability is especially useful when attackers use different methods or new exploits to attack a system.



EXAM TIP Attackers are constantly discovering new exploits, or new ways to exploit vulnerabilities. A honeypot gives administrators and security professionals an opportunity to observe new exploits.

A *honeynet* is an extension of a honeypot and includes two or more servers configured as honeypots. The goal of the honeynet is to simulate network activity similar to what an attacker would see in a live environment. In contrast, a single honeypot may not have any activity, enabling an intelligent attacker to determine that it is a trap.

Padded cells combine intrusion prevention systems with a honeypot. When the system detects an attack, it transfers the attacker to a system similar to a honeypot. In this context, the honeypot is called a padded cell. Just as a honeypot is a system set up with fake data, the padded cell also includes fake data. The data is intended to look useful to keep the interest of the attacker, but it doesn't have any real value. Just as administrators can track and observe the attacker in the honeypot, they can also watch activity within the padded cell.

Using Security Information and Event Management Tools

Security information and event management (SIEM) applications monitor an enterprise for security events. They are very useful in large enterprises that have massive amounts of data and activity to monitor. Other names for these solutions include security information management (SIM), system event management, and security event management (SEM) application.

Consider an organization with over 3,000 servers. When a situation arises on just one of those servers, administrators need to know as quickly as possible. The SIEM application provides constant monitoring and reporting capabilities to provide real-time reporting.



TIP Security information and event management applications are very useful for any enterprise. They have historically been very expensive, preventing smaller organizations from using them. However, open source solutions such as Security Onion are now available to smaller organizations.

Each of these solutions includes a group of tools used to gather information from systems and network devices on the network, analyze the data, and report on items of interest. It's common for a SIEM application to include a NIDS/NIPS component, but it will include more. For example, if an organization must comply with specific laws, a SIEM application could have components to monitor systems for compliance with the laws and report if a system is no longer compliant. These tools within the SIEM application are combined into a single unified solution with several benefits, including the following:

- **Single interface to monitor all activity** Administrators don't need to monitor multiple tools because the SIEM application provides a single interface.
- **Database capabilities** Collected data is stored in a database that can easily be searched and analyzed for items of interest. Data includes logs collected from servers and network devices such as firewalls, routers, and switches. Log types include security logs, system logs, firewall logs, backup logs, and more.

- **Ability to define or fine-tune items of interest** SIEM applications come with predefined alerts, but they can be modified with different threshold levels or different items to search. For example, if management is interested in specific events occurring on a server, such as the use of privileges by anyone with an administrator account, they can have personnel configure the SIEM application to issue an alert each time these events occur.
- **Alerting capabilities** SIEM applications have the ability to issue an alert on any item of interest and provide different types of notifications. They can provide visual indications, such as pop-ups through the application's interface, and send notifications via e-mail or text messages.
- **Secure storage of data** Because a SIEM application is centrally located, the data can be protected from an attacker. In other words, the logs for individual monitored systems are stored remotely with the SIEM application. If any of the systems are attacked, the logs stored with the SIEM application remain intact.

Performing Security Assessment Activities

Many organizations perform regular security assessments. Administrators use them to check the security posture of the organization and discover vulnerabilities. After performing a security assessment, they report their findings and recommendations, if appropriate, to management. Management can approve the recommendations to implement security controls to reduce the identified vulnerabilities, increasing overall security.

The most common type of security assessment is a vulnerability assessment. Organizations may take this a step further and perform a penetration test that attempts to exploit weaknesses discovered by a vulnerability assessment. Penetration tests are more intrusive and have the capability to cause damage if not approached cautiously.



TIP This section covers vulnerability assessments and penetration tests. Chapter 10 covers security audits, which are also a form of a security assessment.

Vulnerability Assessments

Vulnerability assessments attempt to discover vulnerabilities and can take many forms. Some assessments attempt to discover whether an organization is susceptible to social engineering attacks. Other assessments attempt to discover whether the organization has technical weaknesses.

For example, phishing attacks are common, and an organization may want to test its employees' knowledge of phishing to see whether they are likely to respond. An online company, KnowBe4 (www.knowbe4.com), provides a free service that will send a bogus phishing e-mail to employees to determine how many employees are likely to respond to a phishing e-mail. Ideally, zero employees will respond, but if even one employee responds, it could compromise the organization's resources. Tests such as the one provided by KnowBe4 often show that many employees aren't educated about the risks of phishing e-mails and respond or click on links.

APT1 and Phishing

The security firm Mandiant discussed the attack life cycle in its 2013 report titled *APT1: Exposing One of China's Cyber Espionage Units* (previously introduced in Chapter 5). The life cycle typically starts with spear phishing (targeted phishing attacks) and includes subject lines and text within the e-mail that are relevant to the recipient. These e-mails typically include a link to a malicious ZIP file.

If a user opens the ZIP file, it installs a backdoor on the user's system. This backdoor causes the user's system to initiate connections with the attackers' command and control server. Attackers configure the command and control server to respond by sending commands back to the user's system.

One common method they use is to embed commands within a web page as HTML tags. These commands allow the attacker to perform internal reconnaissance on the infected user's network using standard commands. For example, **net view domain** will show the domain name and **net view /domain:domainname** (with the name of the domain in place of *domainname*) will list all of the computers in the domain. The attackers typically dump this list to a text file with a command like this: **net view /domain:domainname >> 1.text**. Later, they transmit the text file back to the command and control server and use this information to install multiple backdoors on other computers in the network.

Think about this for a moment. One employee in an organization opens a ZIP file and the result is a massive infiltration of backdoors installed on multiple systems. Once the attackers establish a foothold, they begin collecting data, compressing it into multiple RAR files, and sending these files to the attackers.

Other security consultants use social engineering tactics to gain knowledge of an organization and then use this knowledge to gain more and more access. For example, an organization may hire a security consultant to see what information the consultant can learn about the organization just from calling employees or going through its trash.

Technical vulnerability assessments attempt to discover technical weaknesses. For example, vulnerability scanners can scan systems and networks to discover existing vulnerabilities.

Test Types

Personnel performing a vulnerability assessment have different levels of knowledge, based on the type of test they are performing. There are three common test types:

- **White box or internal testing** Testers perform the assessment within the organization. They have full access to the internal network and know the network infrastructure, including what systems it hosts. Security professionals who are employees within the organization commonly do internal testing. Other names for this type of testing include clear box testing, glass box testing, and transparent box testing.

- **Black box or external testing** Testers don't have any knowledge of the internal network prior to starting the testing. Consultants hired to test an organization's security vulnerabilities often do black box testing.
- **Gray box testing** This is a cross between white box testing and black box testing. It can be done internally or externally, but the testers have at least some level of knowledge about the network.



TIP These test types are also called full knowledge (white box), zero knowledge (black box), and partial knowledge (gray box) tests.

Figure 8-6 shows how a white box internal test is able to test multiple systems on the internal network, the DMZ, and even firewalls protecting the network. Chapter 5 defines white hats, black hats, and gray hats. These terms are completely different from white box testing, black box testing, and gray box testing. For example, testers working for an organization with no malicious intent are white hat testers. White hat testers can perform white box testing, black box testing, or gray box testing.

It's worth stressing that many attackers use the same tools organizations use to assess an organization's vulnerabilities for potential attacks. In other words, if you can discover a vulnerability with one of these tools, the odds are very good that an attacker can also discover these vulnerabilities.

Vulnerability-Scanning Tools

Many vulnerability-scanning tools are available to automate scans of a network. Some tools are comprehensive and do multiple functions, while other tools perform a single function. These tools attempt to perform overall reconnaissance of a network. The result is a listing of the machines that are operating on the network along with the likely services running on each system.

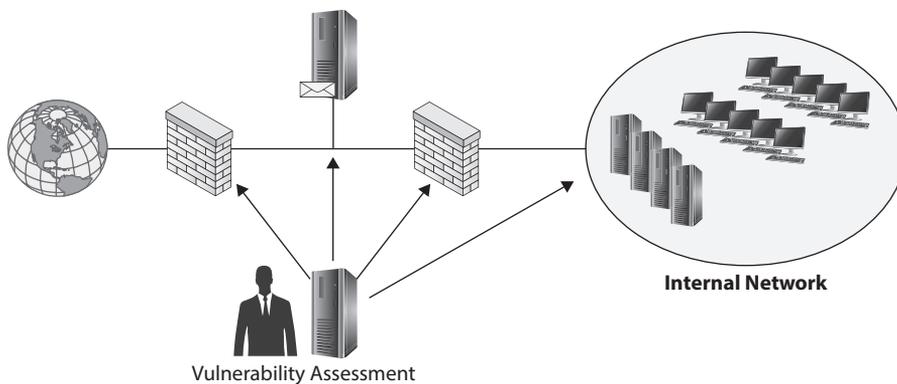


Figure 8-6 White box testing launched internally to check systems

Some of these tools start with an Internet Control Message Protocol (ICMP) sweep to identify computers based on their IP addresses. They follow this with a port scan on each of the computers to identify which ports are open on individual computers. Next, they try to identify more details about the system, typically by analyzing packets sent from the system.

As an example, Table 8-1 shows the values of two observable items in a packet. The recommended default value for the initial Internet Protocol (IP) Time To Live (TTL) is 64, but this is only a recommendation and different operating systems have different values. It refers to the number of routers, or hops, a packet will traverse before it is discarded. Similarly, different operating systems use different TCP window sizes for the initial packets in a TCP session. In some cases, different versions of operating systems have different values. For example, Windows XP uses a TCP window size of 65,535 bytes, but newer Windows versions use a window size of 8,192 bytes.

The initial TTL and TCP window size are only two items in a packet that identify details about a system. Vulnerability scanners use multiple comparisons to identify a system. This process is often referred to as *fingerprinting* the system.



TIP Scanning tools often use sniffers to capture and analyze packets. Chapter 5 discusses sniffers (also called protocol analyzers) in more depth.

Nmap (short for Network Mapper) can locate and document all the devices on a network as part of a reconnaissance scan. Nmap uses scanning methods first to discover all the devices based on their IP addresses, and then uses scans on each individual system to identify services running on each system. Nmap can often identify a system's operating system, services running on the system, the existence of a host-based firewall, and much more.

Nessus is a vulnerability-scanning tool used to discover vulnerabilities in individual systems. It can use Nmap to discover systems on the network during an initial reconnaissance scan. It then uses other methods for the fingerprinting scan. It can perform detailed analysis of individual systems, including auditing passwords for complexity, checking systems for compliance with specific security policies, and checking for the existence of personally identifiable information (PII) on a system.



EXAM TIP Many laws mandate the protection of PII. Tools such as Nessus can be used by organizations as part of a compliance check to ensure that they are in compliance with existing laws. Organizations keep the results of these scans to prove that the organization is actively involved in compliance monitoring.

Operating System	IP Initial TTL (hops)	TCP Window Size (bytes)
Windows XP	128	65,535
Windows (after XP)	128	8,192
Linux	64	5,840
FreeBSD	64	6,5535
Cisco IOS	255	4,128

Table 8-1 TTL and Window Sizes

There are many other vulnerability-scanning tools available. Security professionals testing networks for vulnerabilities can use these tools, but remember that criminals attempting to break into networks can also use them.

No matter which tool you use, you can expect the security professionals to provide detailed reports documenting the findings. Of course, an extremely important step in the vulnerability assessment process is to analyze these results and address any vulnerabilities that can be mitigated. Data in these reports is sensitive and should be protected. If attackers have access to the reports, they can use the information to launch attacks on documented vulnerabilities.

Reviewing Infrastructure Security Configurations

Infrastructure security refers to securing the hardware elements of a network. It's important to review these security elements periodically to verify they are up to date with current security practices.

Administrators performing daily tasks in a network often get used to the configuration and can overlook some things. Because of this, organizations often bring in an outside consultant to examine the infrastructure and provide recommendations. A security consultant will typically perform the following tasks and provide recommendations to improve security:

- **Assess network vulnerabilities** This can be done by running various vulnerability-scanning tools such as Nessus.
- **Verify appropriate host hardening** This check ensures that servers are running only the protocols and services they need and all other protocols and services are disabled. It also ensures that host-based firewalls are enabled and systems are up to date with current patches.
- **Check border security** This includes checking all paths into the network from external sources such as the Internet. The consultant would check network devices such as routers, firewalls, UTM appliances, and the DMZ configuration.

Instead of hiring outside consultants to perform these reviews, some organizations have dedicated IT security staff. If these personnel are given the time and resources to focus on security, they are able to perform these assessments on a routine basis.

Order of Steps

Vulnerability assessments typically use the following order of steps:

1. **Gain permission from management.** The importance of this step cannot be overstated. Senior management from any organization typically interprets unapproved vulnerability scans as attacks. This often results in the loss of a job or even legal action. Most security professionals require written approval before starting an assessment. This is so important that if someone is willing to perform an assessment without written approval, they should be viewed with suspicion.

2. **Discovery.** In this step, personnel perform vulnerability scans on the network to identify weaknesses. It often includes reconnaissance steps to document the environment, followed by fingerprinting scans to identify specifics about the systems, such as the operating system or running services. The detailed scans on individual systems can detect specific vulnerabilities for each of the scanned systems.
3. **Analyze results.** Vulnerability scanners provide reports that personnel must analyze to determine their validity. It's possible for a scanner to misidentify a system or report a false positive. Security professionals who regularly run scans recognize the weaknesses and the strengths of the tools. For example, a scan may regularly report that a firewall is not enabled on a host simply because it doesn't recognize a third-party firewall. However, when new vulnerabilities appear on reports, they deserve extra scrutiny to determine their validity.



NOTE Some scanners can be configured to eliminate false positives for specific events. For example, it may be possible to eliminate a false positive if the scanner doesn't recognize a third-party firewall. The challenge is ensuring that the configuration change doesn't create any false negatives. In this example, you want to ensure the scanner still reports when a system doesn't have any firewall enabled.

4. **Document vulnerabilities.** Next, personnel document the vulnerabilities in a report. Some organizations want this report turned in as soon as possible so that management has an idea of the vulnerabilities. For example, in one large network operations center where I worked, a security group performed detailed scans once a month. As soon as the group was able to compile the results, it published the report, making it available to key management personnel in other departments. In this situation, the security group simply performed the scans and reported the results, and other departments were responsible for identifying and recommending methods to remediate the vulnerabilities.
5. **Identify and recommend methods to reduce the vulnerabilities.** Next, personnel identify methods that can mitigate the vulnerabilities. These could be as simple as recommending a configuration change or as complex as recommending a detailed patch management process to ensure all systems are kept up to date.
6. **Present recommendations to management.** Management has the responsibility to examine recommendations and either approve or reject the recommendations. They have the overall view of the organization and the responsibility for any losses associated with residual risks for any rejected recommendations.
7. **Remediate.** Last, personnel implement any recommendations approved by management. It's worth stating the obvious here: If a fix isn't implemented, the vulnerability remains and attackers can exploit it. If management has approved a recommendation, it should be implemented as soon as possible.



EXAM TIP The last step (and arguably the most important step) in a vulnerability scan is to implement the fixes. Many organizations often repeat the scan after implementing the fix to verify that they have actually removed the vulnerability as expected.

False Positives and the Importance of Results Analysis

In some cases, a vulnerability assessment may indicate that a vulnerability exists even though it doesn't. This is a false positive. For example, a scanning tool may indicate that a server is missing a patch, but investigation verifies that the patch is installed. This is often a weakness in the vulnerability-scanning tool, and some tools are more prone to false positives than other tools are.

However, the important thing to recognize is that false positives can occur. Just because the tool indicates a vulnerability exists, that doesn't mean it does. Administrators or security analysts examine the results to determine their validity.



TIP False positives from vulnerability tools are often unique to each tool. One tool may be able to effectively determine if a patch is installed, while another tool may report erroneously that a patch is missing.

Vulnerability tests can sometimes give false negatives also. As an example, imagine management decides not to implement an operating system patch on a server because

Get Out of Jail Free Card

If an organization hires you to perform a vulnerability assessment or penetration test, it's important that you obtain written permission from management before starting. If not, you could face legal action after doing the test.

As an example, several security professionals were granted access to an organization's campus to secretly perform vulnerability testing. Some administrators detected their activity and did some digging. With the help of internal security personnel, they located the pseudo-attackers in an area of a building that was under renovation.

These pseudo-attackers were dressed as construction workers and working behind some tarps. With all the computer equipment they had running, it was clear they weren't construction workers. However, as soon as they were discovered, they quickly showed their written authorization, calling it their "get out of jail free" card from Monopoly fame. A phone call confirmed that the written authorization was valid and they were not arrested.

In contrast, if they didn't have written authorization, the result would have been much different. They would have been taken to the campus security office and perhaps arrested. They probably wouldn't have spent much time in custody, but what if the individuals who verbally OK'd the test weren't available? Worse, what if the managers who approved the test recalled a different version of what they authorized? The team might be accused of exceeding their authority.

Written documentation avoids all the problems.

it breaks an application. A vulnerability scan should indicate the system is vulnerable because the patch isn't applied, but if it doesn't, it's a false positive. This can give a false sense of security because a vulnerability scan may indicate that a network is secure, yet actual vulnerabilities exist.

Penetration Tests

Penetration testing takes a vulnerability assessment one step further. Instead of stopping after discovering vulnerabilities, a penetration test attempts to exploit the vulnerability. While a penetration test provides verification that an attacker can (or can't) actually exploit a system, it comes with significant risks.

It's important for testers to recognize the potential impact of security testing. It's possible for a penetration test to cause damage resulting in the loss of confidentiality, integrity, or availability of a system or data. This is never a goal, but due to the nature of penetration testing, it is possible, so testers must be cautious when testing.

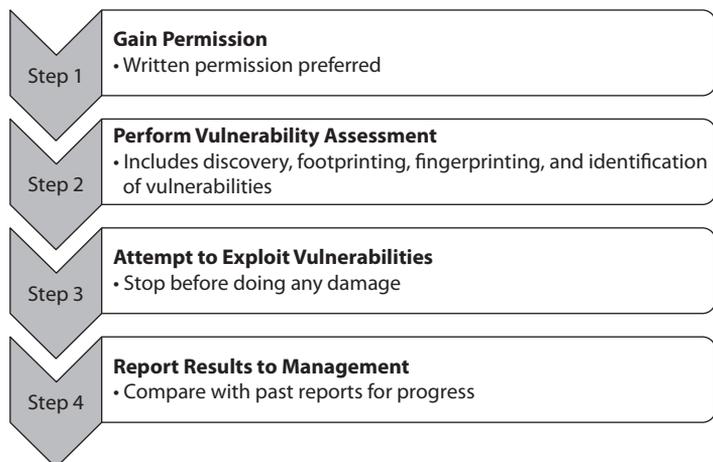
Both internal and external penetration tests are possible, and they are classified as white box (full knowledge), black box (zero knowledge), and gray box (partial knowledge) tests, just as vulnerability tests are classified. The amount of information given to the testers is dependent on the goals of the test.

Figure 8-7 shows the overall steps of a penetration test, and the following sections cover them in more depth.



EXAM TIP A penetration test is an intrusive test. However, all penetration tests should stop before causing damage to a live system or network. The goal is to prove that exploiting a vulnerability can affect the mission of the organization. The test should stop before it actually affects the mission of the organization.

Figure 8-7
Penetration
test steps



Gain Permission

Gaining permission is an even more important step in a penetration test than it is with a vulnerability test because it's possible that a penetration test can actually result in a loss. The permission needs to be in writing and should identify the goals of the test.

Perform Vulnerability Assessment

In the vulnerability assessment stage, you go through similar steps as you would in a typical vulnerability assessment. While the actual steps of the vulnerability assessment can vary depending on the scope, it's very likely it will include the following components:

- **Discovery or reconnaissance** In this stage, you collect information about the targets using any available means. This could use tools such as Nmap to complete a reconnaissance scan, or use existing documentation if the test is carried out as a white box test. Black box testers often need to use a variety of social engineering tactics to gain information about the organization and the network.
- **Fingerprinting** Once the network is mapped, individual systems are identified and probed to discover details about them. This includes the operating system, service packs, patches, open ports, running services, and running applications.
- **Identify vulnerabilities** The tester then analyzes the collected data to determine what vulnerabilities exist and can be exploited.

If this were a vulnerability test, the data would be put into a report and the test would be complete. However, a penetration test adds the next step: an attempt to exploit one or more vulnerabilities.

Attempt to Exploit Vulnerabilities

In this step, the tester attempts to use methods similar to what an attacker would use to exploit the vulnerabilities. At this stage, it's very important to understand the scope and purpose of the test. Most organizations want to know what can be exploited, but not at the expense of an outage. In other words, the tester should take the test as far as possible without causing any damage.

Of course, the danger with this strategy is that you don't always know how far is too far until you get there. Testers should be acutely aware of the potential damage that an exploit can cause and be very cautious when the exploit can be damaging.

Some exploits result in the disclosure of data. In this case, testers can gather as much data as possible without causing harm. The testers can show the data they've collected to demonstrate that an attacker can also access it. Other exploits can modify or corrupt data, and these types of tests should not be performed on live production servers. However, it is possible to create a test server configured identically to the production server. Testers then attempt the penetration test on the test server. If the test on the test server succeeds, it provides proof that the exploit is possible without affecting actual resources.

Report Results to Management

The results of the penetration test are compiled into a report, along with recommendations for controls or safeguards to mitigate any of the vulnerabilities. Just as with a vulnerability report, management decides on what risks to mitigate and what risks to accept.

Past reports are valuable to compare. For example, a past report may have documented a vulnerability, and management may have approved a control to mitigate the risk. If the vulnerability still exists, it deserves some investigation. Was the control implemented? Was it completely ineffective, or did it actually reduce the vulnerability? Has the threat technique been modified to circumvent the original control? Answers to these questions help determine the next steps to take.

Chapter Review

An intrusion detection system (IDS) provides continuous monitoring and sends alerts to report suspicious activity. Administrators investigate alerts to determine whether they are actual attacks or false positives. NIDS agents are installed on network devices and send data back to a central network-based IDS (NIDS) management server. This server analyzes the activity throughout the network and provides notification when an alert occurs. A host-based IDS (HIDS) monitors activity on an individual system, such as a server or workstation, and provides notifications as alerts. Notifications are typically recorded in logs but can also be sent as e-mails or text messages. An intrusion prevention system (IPS) is placed inline with the traffic and can modify the environment to block attacks in progress.

IDSs and IPSs use two primary detection methods. A signature-based (or knowledge-based) system has a database of known attacks and compares activity with this data. This is similar to how AV software uses signature files to detect malware. An anomaly-based (or behavior-based) system starts with a baseline of normal activity. It then monitors ongoing activity to determine whether there is a significant change or anomaly. Anomalies are reported as alerts and need to be investigated. Many systems use a hybrid model of both signature-based and anomaly-based detection methods.

One of the benefits of NIDSs over HIDSs is that the log files are stored remotely on the central server. A HIDS stores the logs on the host, and an attacker can modify or delete the logs, erasing the record of the attack.

Many times an organization wants to detect unauthorized changes. Tools such as file integrity checkers monitor files and can detect unauthorized modifications.

Honeypots are systems with fake data designed to attract attackers, diverting them from a live network. They also give security professionals an opportunity to observe attacks in action. Honeynets are groups of two or more honeypots to simulate a network. Padded cells combine an IPS with a honeypot by transferring an attacker into the honeypot after detecting an attack.

Security information and event management (SIEM) applications are unified monitoring systems used by large enterprises. They combine the strengths of multiple tools into a single management tool.

Most organizations perform security assessments to detect potential security issues. A vulnerability test attempts to discover vulnerabilities and provide recommendations to mitigate the associated risks. Vulnerability tests can be white box (full knowledge), black box (zero knowledge), or gray box (partial knowledge) tests. Testers can use nontechnical means, such as social engineering tactics, and use technical means with tools such as Nmap and Nessus.

Penetration tests go a step further than a vulnerability test. They attempt to exploit a vulnerability. Penetration tests are more intrusive and have the potential to cause damage, so they should be performed with utmost caution and stopped before inflicting any actual harm. Testers should obtain written permission prior to beginning either a vulnerability assessment or a penetration test. Both vulnerability assessments and penetration tests provide a report to management as the last step.

Questions

1. You want to monitor the network for possible intrusions or attacks and report on any activity. What would you use?
 - A. HIPS
 - B. HIDS
 - C. NIDS
 - D. AV software
2. You want to monitor a server for potential attacks. Of the following choices, what is the best choice?
 - A. HIDS
 - B. NIDS
 - C. Anomaly-based IDS
 - D. Signature-based IDS
3. Which of the following choices identifies a major drawback associated with a host-based IDS (HIDS)?
 - A. It is very processor intensive and can affect the computer's performance.
 - B. The signatures must be updated frequently.
 - C. It does not support anomaly-based detection.
 - D. It stores the logs on remote systems.
4. Of the following choices, what best describes an IPS?
 - A. An active antivirus program that can detect malware
 - B. An inline monitoring system that can perform penetration testing
 - C. An inline monitoring system that can perform vulnerability assessments
 - D. An inline monitoring system that can modify the environment to block an attack

5. What type of control is a NIDS?
 - A. Corrective
 - B. Detective
 - C. Deterrent
 - D. Preventive
6. Which of the following identifies a system that requires a database to detect attacks?
 - A. Anomaly-based IDS
 - B. Signature-based IDS
 - C. HIPS
 - D. NIPS
7. You have recently modified the network infrastructure within your network. What should be re-created to ensure that the anomaly-based NIDS continues to work properly?
 - A. Signature database file
 - B. Baseline
 - C. Router gateways
 - D. Firewalls
8. How does a behavior-based IDS detect attacks?
 - A. It compares current activity against a baseline.
 - B. It compares current activity against a database of known attack methods.
 - C. It compares current activity with antivirus signatures.
 - D. It monitors activity on firewalls.
9. Of the following choices, what best describes a whitelist as a replacement for a HIDS?
 - A. A listing of websites that a user can visit, blocking access to all other websites
 - B. A listing of applications that a user can run, blocking attempts to run any other applications
 - C. A listing of MAC addresses blocked through a firewall, allowing traffic from all other systems
 - D. A listing of suitable vendors for IPSs
10. What logs are most valuable after an attack?
 - A. Logs on a remote system
 - B. Logs on local systems that have been attacked
 - C. Logs for local firewalls
 - D. Logs for antivirus events

11. Which of the following can detect if a system file has been modified?
 - A. Encryption algorithm
 - B. Anomaly-based detection
 - C. Signature-based detection
 - D. File integrity checker
12. Of the following choices, what is a primary purpose of a honeypot?
 - A. To give administrators an opportunity to observe new exploits
 - B. To give administrators an opportunity to observe new controls
 - C. To give administrators an opportunity to perform vulnerability tests
 - D. To give administrators an opportunity to perform penetration tests
13. An external organization is performing a vulnerability test for a company. Officials from the company give this group some information on the company's network prior to the test. What type of test is this?
 - A. White box test
 - B. Gray box test
 - C. Black box test
 - D. Internal test
14. Your organization has contracted with a security organization to test your network's vulnerability. The security organization is not given access to any internal information from the company. What type of test will the organization perform?
 - A. White box testing
 - B. Gray box testing
 - C. Black box testing
 - D. Partial knowledge testing
15. What is an important first step in a vulnerability assessment?
 - A. Document vulnerabilities
 - B. Fingerprinting
 - C. Reconnaissance
 - D. Gaining approval
16. What is the last step in a vulnerability assessment?
 - A. Discovery
 - B. Analysis
 - C. Remediation
 - D. Document vulnerabilities

17. How does a vulnerability scanner fingerprint a system?
- A. With a biometric scanner
 - B. Using an ICMP sweep
 - C. By identifying its IP address
 - D. By analyzing packets
18. A vulnerability assessment reports that a patch is not installed on a system, but you've verified that the patch is installed. What is this called?
- A. Anomaly-based vulnerability
 - B. Signature-based vulnerability
 - C. False negative
 - D. False positive
19. What's the primary difference between a penetration test and a vulnerability assessment?
- A. A vulnerability assessment includes a penetration test, but a penetration test does not include a vulnerability assessment.
 - B. A penetration test is intrusive and can cause damage, while a vulnerability assessment is passive.
 - C. A vulnerability assessment is intrusive and can cause damage, while a penetration test is passive.
 - D. They are basically the same, but with different names.
20. When should a penetration test stop?
- A. After discovering the vulnerabilities
 - B. After discovering the threats
 - C. Before causing any damage
 - D. Before discovering the exploits

Answers

1. **C.** A network-based IDS (NIDS) monitors the network in real time and sends alerts to report suspicious activity. A host-based intrusion detection or prevention system (HIDS or HIPS) cannot monitor network activity. AV software detects malware, but not necessarily network attacks.
2. **A.** A host-based IDS (HIDS) can monitor a single computer (such as a server) for possible attacks and intrusions. A network-based IDS (NIDS) monitors network activity. Both HIDS and NIDS can use either anomaly-based detection or signature-based detection.

3. **A.** A HIDS is very processor intensive and can negatively affect the computer's performance. All IDSs using signature-based detection need to have their signatures updated frequently. A HIDS will typically support both signature-based detection and anomaly-based detection. Another drawback of a HIDS is that it stores the logs on the local computer, and an attacker might be able to delete the logs.
4. **D.** An intrusion prevention system is an inline monitoring and detection system that can modify the environment (such as by changing ACLs or closing half-open connections) to block an attack. Although it may be able to detect some malware such as worms, this isn't the best definition. It is not used for vulnerability assessments or penetration tests.
5. **B.** A network-based IDS (NIDS) is a detective control, as it detects potential attacks as they are occurring. An active IDS is a corrective control because it can take action to reverse the effects of an attack by changing the environment, but all IDSs are not active. A deterrent control attempts to deter would-be attackers from attempting an attack, but attackers don't know if a network has an IDS. An intrusion prevention system (IPS) is a preventive control because it can prevent an attack from reaching a network, but an IDS does not prevent the attack.
6. **B.** A signature-based IDS compares activity against a signature file (or database of signatures) to identify attacks. An anomaly-based IDS requires a baseline. Both a host-based IPS (HIPS) and a network-based IPS (NIPS) can use either anomaly-based or signature-based detection methods.
7. **B.** An anomaly-based NIDS compares current activity against a baseline to determine abnormal behavior, and this baseline should be updated when the network is modified. This provides the NIDS with an accurate baseline. A signature-based NIDS uses a signature database file. Router gateways and firewalls do not need to be re-created for the NIDS.
8. **A.** A behavior-based IDS detects attacks by comparing current activity against a baseline. A signature-based IDS detects attacks by comparing network activity with a database of known attack methods. An IDS does not use antivirus signatures. While an IDS monitors activity on firewalls, this doesn't identify how it detects attacks.
9. **B.** A whitelist can include a list of applications that a user can run and block all other applications. Some security professionals are suggesting this as a replacement for both HIDSs and AV software. A proxy server can block access to specific websites, but this isn't a replacement for a HIDS. A MAC address whitelist identifies addresses that are allowed, not blocked. A whitelist is not a list of vendors.
10. **A.** Logs held on remote systems are the most valuable because an attacker is less likely to have modified them. Attackers can modify logs on a local system (including local firewall logs) during an attack. Although AV logs may be useful in attacks involving malware, they aren't very useful if attackers didn't use malware during the attack.

11. **D.** File integrity checkers use hashing techniques to detect whether modifications of system or other important files have occurred. File integrity checkers use hashing algorithms, not encryption algorithms. Intrusion detection and prevention systems use signature-based detection to detect known attacks and use anomaly-based detection to detect previously unknown attacks, but these are not related to detecting modifications to system files.
12. **A.** A honeypot entices would-be attackers, luring them away from the live network, and gives administrators an opportunity to observe the attacker. The honeypot is a security control, but it doesn't provide an opportunity to observe new controls. It does not perform vulnerability or penetration tests.
13. **B.** A gray box test (also called a partial knowledge test) is performed with some internal knowledge and can be performed either internally or externally. A white box test (also called a full knowledge test) is performed with full access to internal documentation. A black box test (also called a zero knowledge test) is performed without any inside knowledge of the organization. It is not an internal test because an external organization is performing it.
14. **C.** A black box test (also called a zero knowledge test) is performed without any inside knowledge of the organization. A white box test (also called a full knowledge test) is performed with full access to internal documentation. A gray box test (also called a partial knowledge test) is performed with some internal knowledge.
15. **D.** An important first step in a vulnerability assessment is gaining written approval from management. Discovery includes a reconnaissance scan and fingerprinting techniques, but should not be done without approval. Documenting the vulnerabilities comes after discovery and analysis.
16. **C.** The last step in a vulnerability assessment is remediation of vulnerabilities using controls approved by management. Discovery, analysis, and documentation all occur after gaining approval from management, but before remediation.
17. **D.** A vulnerability scanner fingerprints a system by analyzing packets sent out by the system. In this context, fingerprinting is a metaphor and it doesn't use a biometric scanner. An Internet Control Message Protocol (ICMP) sweep identifies IP addresses on a network, but the IP addresses do not fingerprint a system.
18. **D.** A false positive occurs when a vulnerability assessment tool indicates that a vulnerability exists when it actually does not exist. A false negative occurs when a vulnerability assessment tool indicates that a vulnerability doesn't exist when it actually does exist. Anomaly-based and signature-based vulnerabilities are detection methods of IDSs and are not associated with vulnerability assessment tools.
19. **B.** A penetration test is intrusive; it includes a vulnerability assessment, attempts to exploit discovered vulnerabilities, and can cause damage. Vulnerability assessments are not intrusive and do not cause damage.
20. **C.** Most penetration steps stop before executing an exploit that can cause damage. Penetration steps include the basic vulnerability assessment components of identifying vulnerabilities, threats, and exploits, and then follow these steps with an attempt to exploit a discovered vulnerability.

Controls and Countermeasures

In this chapter, you will learn about

- Goals of controls and countermeasures
- Differences between deterrent, preventive, detective, corrective, and compensating controls
- Differences between management, technical, operational, and physical controls
- Differences between policies, standards, procedures, and guidelines
- Basic controls such as change control, configuration management, and patching systems
- Mobile device management methods and other endpoint device security concepts
- Use of RAID subsystems to provide fault tolerance
- Use of failover of clusters to protect against the failure of a server
- Different types of backups, such as full, differential, and incremental

Using Controls, Safeguards, and Countermeasures

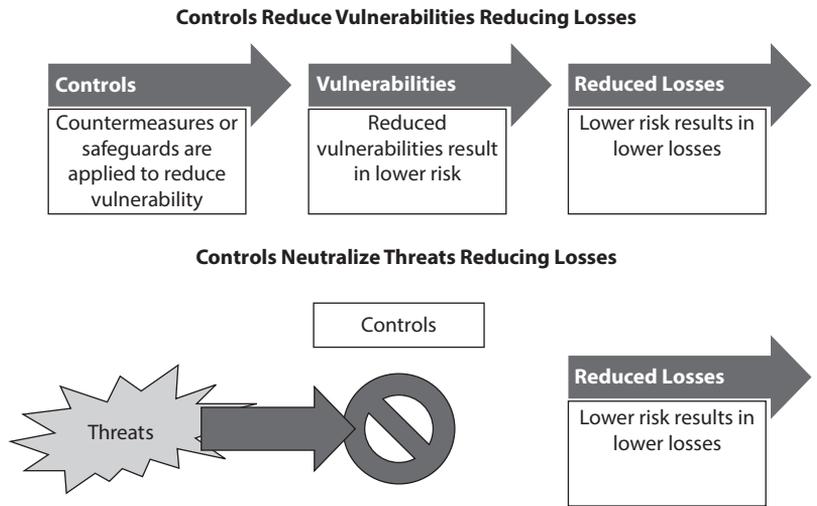
Chapter 7 provided information on risk and explained that risk is the probability or likelihood that a threat will exploit a vulnerability and cause a loss. Chapter 8 presented information on performing vulnerability assessments to detect vulnerabilities. Both chapters mentioned that an organization mitigates risk by implementing security controls. This chapter digs deeper into security controls.

Security professionals often use the terms *controls*, *safeguards*, and *countermeasures* interchangeably. In essence, they are means, methods, actions, techniques, processes, procedures, or devices that reduce the vulnerability of a system or reduce the probability that a threat will exploit a vulnerability and cause a loss. In this chapter, I've primarily used the term *controls*. However, if you see the terms *safeguards* or *countermeasures*, it's important to recognize they mean the same thing.

Security controls help prevent, detect, and correct problems related to security incidents. In general, a security incident is any event that results in a negative impact

Figure 9-1

Controls reduce vulnerabilities, resulting in reduced losses.



on the confidentiality, integrity, or availability of an organization's assets. This includes losses from data breaches, denial of service (DoS) attacks, malicious software (malware) infections, power loss, and failure of environmental systems such as heating, ventilation, and air conditioning (HVAC) systems.

Figure 9-1 shows the two ways that controls are used. They can either reduce vulnerabilities to reduce losses from a risk or attempt to neutralize a threat to reduce losses from a risk. For example, antivirus (AV) software protects against malware. When administrators install AV software and keep it up to date, the system is less vulnerable to malware attacks. The AV software also helps to neutralize the threat of malware.

Controls can be either technical or nontechnical. Technical controls use technical means to reduce risk and are merged with hardware, software, and firmware to either reduce vulnerabilities or reduce the impact of threats. Nontechnical controls are management and operational controls and include items such as written security policies, physical security controls, operational procedures, and personnel training.

For example, an intrusion prevention system (IPS), covered in Chapter 8, is an example of a technical control that attempts to detect and block attacks. User training is an example of a nontechnical control that attempts to educate users and encourage them to avoid risky behavior such as clicking a link within a phishing e-mail.

Performing a Cost-Benefit Analysis

Organizations have a finite amount of money, so they must evaluate controls and purchase controls they perceive as having the best value. A cost-benefit analysis (CBA) attempts to determine the value of a control to the organization. A CBA compares the cost of a control with the potential benefits of the control to help determine if the cost of

the control is justified. The CBA also compares the cost of implementing a control with the cost associated with losses if the control is not implemented.

Personnel often perform a CBA for each recommend control. This makes it easier for management to make an educated decision about whether or not to purchase the control. Chapter 7 presents both quantitative analysis and qualitative analysis methods. A quantitative analysis uses monetary amounts that work well with a CBA.

When evaluating any control, it's important to evaluate both its expected effectiveness in mitigating a risk and its cost. Further, when evaluating the cost of the control, you must consider more than just the initial cost. Controls have several costs that management must consider before approving the purchase. These additional costs include the following:

- **Initial cost** The cost of the product. It could be a one-time purchase price or could entail monthly or annual recurring costs. For example, AV software has both an initial cost and a subscription cost to update definitions regularly. Of course, not using any type of AV software results in recurring costs too.
- **Implementation costs** Costs required to implement the control. Technical controls in particular often require costs to implement them. Personnel may need to test the control for compatibility, resulting in labor costs. A control may require a vendor to install it initially, or may require training for administrators before the administrators can install and implement the control.
- **Compatibility costs** The costs associated with the usability of a system after implementing a control. Although security often requires a balance with usability, it's easy for security professionals to get overzealous and implement controls without considering the impact on usability. If the control affects productivity, it can affect the mission of the organization.



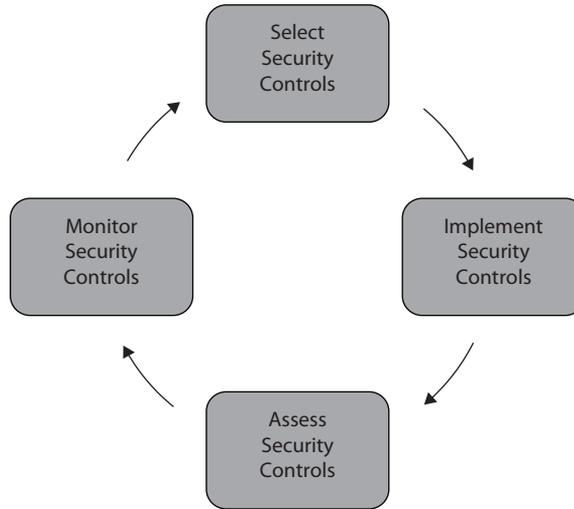
EXAM TIP The cost of a control is justified if the cost of the control is significantly lower than the annual loss expectancy (ALE) without the control. The cost is usually not justified if the cost of the control is significantly higher than the ALE without the control. When the costs and savings are about the same, management may request a return on investment (ROI) analysis to determine whether the cost is justified. The ROI looks at the cost over a longer period of time.

Security Controls Life Cycle

Security controls go through a natural life cycle. This includes selecting, implementing, assessing, and monitoring security controls. This is an ongoing process, and when monitoring indicates that a control is not effective, personnel must evaluate and select additional, or different, controls.

Figure 9-2 shows the overall steps of selecting, implementing, assessing, and monitoring security controls. These steps are derived from National Institute of Standards and

Figure 9-2
Security controls
life cycle



Technology (NIST) Special Publication (SP) 800-53, *Security and Privacy Controls for Federal Information Systems and Organizations*. The steps are as follows:

- **Step 1: Select security controls** After identifying the information system or process that an organization wants to protect, security personnel select and evaluate one or more security controls. The evaluation often includes a risk assessment, including either a quantitative analysis or qualitative analysis, as described in Chapter 7, and a cost-benefit analysis to determine the benefit of the control. While security and administrative personnel recommend security controls, management ultimately selects the security controls.
- **Step 2: Implement security controls** Next, personnel implement the security controls. This can be relatively easy to do or extremely complex, depending on the control. For example, creating a failover cluster to protect availability of a server requires research to purchase the correct hardware, training for personnel to install and maintain the hardware, and a phased implementation plan to put the failover cluster into service.
- **Step 3: Assess security controls** After implementing the security controls, personnel evaluate them to verify they are implemented correctly, operating as expected, and meeting the intended security requirements.
- **Step 4: Monitor security controls** Personnel monitor the security controls on a continuing basis to determine their effectiveness. They also monitor additional factors such as changes to the protected system and changes to external laws, regulations, policies, and guidelines.

It's worth stressing that the life cycle of security controls is an ongoing process. Because threats are constantly changing, security personnel need to constantly monitor the environment for new threats and recommend security controls to mitigate associated risks.

Understanding Control Goals

Controls and countermeasures are typically identified as one of the following three types: preventive, detective, or corrective. A preventive control attempts to prevent losses before they occur, a detective control detects violations, and a corrective control attempts to reverse the impact from a security incident.

Losses to confidentiality, integrity, or availability (CIA) can affect the organization's mission. Figure 9-3 emphasizes that controls attempt to prevent, detect, and correct losses to any of these elements of the security triad.

Some controls combine all three objectives. For example, AV software running on an e-mail server can strip off malicious attachments to prevent them from reaching the user. AV software running on a user's system can detect malware, such as when a user inserts an infected Universal Serial Bus (USB) flash drive into a system. This AV software can also correct the problem by removing the malware from the infected USB flash drive.



EXAM TIP Controls have three primary goals: prevention, detection, and correction. Some controls have only a single goal (prevention, detection, or correction), while other controls attempt to achieve multiple goals.

Preventive

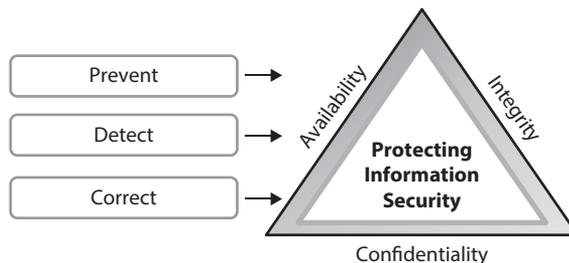
Preventive controls focus on preventing losses due to risks. The most basic type of preventive controls consists of written security policies, standards, and procedures. For example, many of the basic controls, such as separation of duties, least privilege, and strong authentication, are preventive controls identified in written security policies and implemented with different types of controls. Some security controls, such as fences, security cameras, and guards, can deter threats, so can be classified as both preventive and deterrent controls.



NOTE Most documentation uses the term *preventive*. However, the SSCP CIB (and the CISSP CIB) uses the term *preventative*. Both terms mean the same thing in this context.

Figure 9-3

Controls prevent, detect, and/or correct losses to CIA.



The following are some examples of preventive controls:

- Written policies and procedures
- Background checks on prospective employees
- Implementation of separation of duties and least privilege policies
- Strong access control processes (starting with strong authentication)
- Technical password policies that force users to create strong passwords and change them periodically
- Termination processes that ensure personnel disable accounts for terminated employees
- Classification of data (such as public, private, and proprietary) and implementation of varying levels of protection based on the classification
- Encryption of data (both at rest and in transit)
- Security cameras, fences, and guards



TIP Separation of duties can also be detective in nature if it includes a reconciliation function. For example, if accounts payable is separated from accounts receivable, and a third party periodically reconciles the accounts payable and accounts receivable data, the reconciliation process is a detective control that can detect potential anomalies.

Detective

Detective controls identify the event either as it is occurring or after it has occurred. Such controls provide evidence that preventive controls are working (or not working). Information can be used as evidence that a loss has occurred and can be used to provide an immediate response by combining the detective control with a corrective control. The goal is to detect incidents as soon as possible so that either a technical control or personnel can respond to resolve the problem.

The following are some examples of detective controls:

- Forensics analysis
- Physical inventories
- Intrusion detection systems
- AV software that can detect malware installed on a system
- Audit logs (including logs on servers and network resources such as firewalls)
- Reconciliation (comparing different sets of data with each other, such as comparing a database inventory with a physical inventory)



NOTE Many of the controls mentioned throughout this chapter are covered in greater depth elsewhere in this book. I've pointed out the chapters for some of the controls in Notes or Tips. However, you can also use the Table of Contents and the Index to locate where I've covered any of these controls in more depth.

Corrective

A corrective control takes action to reverse the effects of an event and/or restore a disabled or failed control. Many corrective controls often work in conjunction with detective controls. After detecting the event, the corrective control goes into action. For example, after updating signatures on AV software, the AV software might detect an active virus and quarantine it. This reverses the infection. Similarly, a host-based intrusion detection system (HIDS) might detect that the host firewall is disabled and automatically enable it to correct the problem. The following are some examples of corrective controls:

- Intrusion detection systems
- AV software that can remove or isolate malware
- Procedures to back up data and restore backups
- Disaster recovery and business continuity plans



EXAM TIP Many controls combine preventive, detective, and corrective measures. For example, AV software is detective when it identifies malware and corrective when it removes it. Similarly, a backup policy is a preventive control because it attempts to prevent the loss of data. The procedures to restore data from backups are corrective controls because they provide the means to restore data after a data loss.

Other Controls

Beyond preventive, detective, and corrective controls, you should also know about some additional controls, including compensating, deterrent, directive, and recovery controls. Interestingly, the SSCP CIB did not include any of these in past versions, but the current version includes compensating and deterrent controls. It could be argued that each of the following controls could be classified as preventive, detective, or corrective, but these identifications provide a more focused classification:

- **Compensating controls** These controls are in place in case a primary control fails or is unavailable. For example, imagine a small organization employs the principle of separation of duties to ensure that no one controls an entire process in the accounting department. However, occasional illnesses or vacations require one person to do all the work. To compensate for this the organization can implement routine reconciliation audits as a compensating control. The audit will review and validate activity and can detect potential anomalies. You can also

think of these as alternative controls that an organization implements when it isn't feasible or desirable to use the primary control.

- **Deterrent controls** These controls attempt to deter would-be attackers from attempting an attack or even deter users from trying to circumvent policies. For example, just the presence of a guard can deter an attacker from attempting to tailgate into an organization. Similarly, when a proxy server blocks access to a site from within an organization, it could display a web page reminding users of the acceptable use policy (AUP) and informing them that the proxy server records all their online activity. This can deter users from purposely accessing unauthorized pages. Note that a deterrent control is similar to a preventive control. The primary difference is that a deterrent control encourages someone to *decide* not to take a specific action. In contrast, a preventive control doesn't rely on someone doing the right thing, but instead implements methods to prevent them from taking specific actions.
- **Directive controls** These are controls mandated by a higher authority. For example, the Health Insurance Portability and Accountability Act (HIPAA) directs organizations to implement specific safeguards to protect private health information (PHI). The controls that an organization implements to protect PHI are directive controls.
- **Recovery controls** These controls are focused on recovering a system or mission after an outage, such as after a disaster. For example, reliable backups, along with detailed procedures on how to restore backups, are recovery controls. These are often considered corrective controls.



EXAM TIP When preparing for the SSCP exam, focus primarily on the preventive, detective, corrective, deterrent, and compensating controls. However, if you plan to pursue the CISSP certification, you should also have a good understanding of directive and recovery controls.

Thousands of Controls

NIST SP 800-53 (*Security and Privacy Controls for Federal Information Systems and Organizations*) provides extensive coverage of controls. The first 44 pages of this 462-page document provide a solid foundation of the fundamentals related to security controls and the process of selecting security controls. The ten appendixes provide additional details. Appendix F, the “Security Control Catalog,” includes over 200 pages describing specific controls in several different categories.

If you're interested in pursuing a career in IT security, I strongly recommend you download and peruse this document. You can get it from here: <http://csrc.nist.gov/publications/PubsSPs.html>.

One of the most important messages that you can glean from NIST SP 800-53 is that IT security has a lot of depth. There's a lot to know and understand. You aren't expected to know all of the material in SP 800-53 for the SSCP exam, but as you move forward in your IT security career, you'll be much more valuable to employers if you understand these concepts in greater depth.

Comparing the Classes of Controls

In addition to understanding security control goals (such as preventive or detective), you should have an understanding of how controls are implemented. The three primary methods are management, technical, and operational. Figure 9-4 provides an overview, and the following sections explain these classifications in more depth. While it's not a separate class of controls, the SSCP Candidate Information Bulletin (CIB) stresses physical security controls, so this section also includes a "Physical Security Controls" subsection.



TIP Other documents sometimes classify the controls as administrative, technical (or logical), and physical.

NIST SP 800-53 organizes controls into 18 families. In past versions, NIST attempted to match each of these families with specific control classes (such as management, technical, and operational). However, many controls apply to more than one control class, so the current version has dropped that practice. Table 9-1 lists the 18 families. SP 800-53 lists and describes multiple controls within each of these families.

Management/Administrative Security Controls

Management and administrative controls focus on the management of risk and the management of IT security. They include managerial practices that an organization uses to implement security. The following list shows some examples:

- Written documents such as security policies and procedures that provide direction to employees
- Tests and assessments such as risk assessments and vulnerability assessments



Figure 9-4 Control classes: management, technical, and operational

Family	Family
Access Control	Media Protection
Awareness and Training	Physical and Environmental Protection
Audit and Accountability	Planning
Security Assessment and Authorization	Personnel Security
Configuration Management	Risk Assessment
Contingency Planning	System and Services Acquisition
Identification and Authentication	System and Communications Protection
Incident Response	System and Information Integrity
Maintenance	Program Management

Table 9-1 Control Families

- Plans (such as incident response plans) that help an organization respond to incidents
- Screening prospective employees with background checks
- User awareness and training programs to help users take an active role in security



NOTE Chapter 7 covers risk assessments and Chapter 8 covers vulnerability assessments.

Technical Security Controls

Technical (or logical) controls are those that are implemented through technical methods such as hardware, software, or firmware components. The following list shows some examples:

- Identification and authentication controls that allow users to prove their identity
- Access controls that ensure only authorized entities have access to systems and data
- Auditing and accountability controls that track activity and detect unauthorized access
- Network access controls that restrict access to networks to clients that do not meet minimum security requirements
- Encryption protocols that scramble data and provide confidentiality

Operational Security Controls

People implement operational controls on a day-to-day basis to support an organization's security policy during normal daily operations. The following list shows some examples:

- Backup procedures that create reliable data backups
- Routine procedures that ensure personnel safety, such as verifying proper operation of fire prevention and detection equipment

- Routine procedures that ensure system reliability, such as patch management programs
- Configuration and change management practices to increase availability



EXAM TIP Controls and countermeasures can be classified as management, technical, or operational. Management security controls implement administrative practices, often with written documents. Technical controls use technical means within computer systems to reduce risk. People implement operational controls during day-to-day operations.

Physical Security Controls

Physical security controls refer to the controls that you can touch. Many documents group them together with operational controls, and that is accurate. However, within the context of security, physical security controls are so important that it's worth listing them separately. NIST SP 800-53 identifies a family of Physical and Environmental Protection controls that falls into this class of controls. These include any controls that control, restrict, and/or monitor physical access and provide environmental protection (such as temperature and humidity controls). The following list shows some examples:

- Perimeter barriers such as fences and bollards outside the building and locked doors inside the building
- Lighting to illuminate potential areas where intruders may try to break into a building
- Access badges that personnel wear while inside secure areas
- Primary and backup power sources, such as generators and uninterruptible power systems
- Heating, ventilation, and air conditioning systems that control temperature and humidity

When planning physical security controls, it's important to consider personnel safety and always prioritize personnel safety over IT security. As an example, electronic door locks require power to open. If a building loses power, exit doors using electronic door locks should fail in such a way that the doors unlock. This allows personnel within the building to exit. In contrast, if the doors failed in such a way that they remained locked, a fire that affected power could trap personnel inside the room.

Combining Control Goals and Classes

When taking security tests such as the SSCP exam, you're often required to know both the control goal (such as preventive, detective, or corrective) and the control class (such as management, technical, operational, or physical). For example, a locked door to a room is a preventive, physical security control. It is a preventive control because it prevents

individuals from going into the room, and it is a physical security control because you can touch the locked door.

Think about an audit log that security professionals use to identify security violations. How would you classify it? It is a detective, technical security control. It is a detective control because security professionals use it to detect security violations after they've occurred and been recorded in the log. It is a technical control because it uses technical means to record events.

It's impossible to list all the possible controls along with their combinations in this chapter. However, if you understand the control goals (such as preventive, detective, and corrective) and you understand the control, you'll find that it is relatively easy to match the two.

As an example, what is the goal of bollards placed by the entrance to a building? Bollards are vertical metal or cement poles. They act as barricades to prevent someone from driving a vehicle through the entrance, so they are preventive controls.

Similarly, if you understand the control classes (such as management, technical, operational, and physical) and you understand the control, you'll find that it is relatively easy to match the two. Bollards are a physical security control because they use physical methods (something you can touch) for security. In other words, bollards are preventive, physical security controls.

Exploring Some Basic Controls

There are thousands of different security controls, and most (if not all) are documented in NIST SP 800-53. Obviously, this chapter cannot include all of the controls covered in NIST SP 800-53. However, the following sections introduce some common controls that you may see on the SSCP exam.

Hardening Systems

Hardening a system is the practice of making it more secure from its default configuration. Administrators harden systems by following basic security practices such as those in the following list:

- **Remove or disable unused protocols** If a protocol is not being used, it should not be installed. Many operating systems include default protocols. Administrators should verify whether a system needs all installed and enabled protocols, and then either remove or disable any unused protocols. Every installed protocol increases the attack surface of the system. In contrast, by removing unused protocols, administrators decrease the attack surface. Attackers cannot attack a protocol that is not installed and running on a system.
- **Remove or disable unneeded services** This is similar to removing or disabling unused protocols. Many systems include and enable default services. However, not all systems need these default services. Administrators should verify the services that are needed, and disable unneeded services. Attackers cannot attack a

disabled service. In some cases, administrators disable protocols by disabling the related service. However, if malware can get administrative access to a system, it can enable the protocol by starting the service. As an example, some malware has used Trivial File Transfer Protocol (TFTP) to infect other systems on a network. If the related TFTP service was disabled, the malware simply enabled the service, which enabled the protocol.



EXAM TIP You can reduce the attack surface of a system by removing or disabling unused protocols and services. With these protocols and services removed or disabled, an attacker has fewer opportunities to attack a system.

- **Change defaults** If the system starts with any default configurations that represent a risk, change them before putting the system into service. The classic example is default passwords. For example, wireless routers include a default password that, if not changed, attackers can use to easily access the device and modify key settings.
- **Keep systems up to date** Both operating system software and application software develop bugs. Attackers can exploit many of these bugs, so it's important to keep systems up to date. Vendors regularly release updates, patches, and fixes to resolve these bugs. These updates should be tested and applied as soon as possible.
- **Enable firewalls** Many systems have built-in firewalls, and many security suites provided by AV software companies include third-party firewalls. Enabling host-based firewalls reduces the possibility of successful attacks.
- **Include AV software** Threats from malware are constant today, and unprotected systems will quickly become infected. AV software protects systems from a wide variety of malware threats. It's also important to keep the AV software signatures up to date.

Policies, Standards, Procedures, and Guidelines

Policies are high-level documents used to provide guidance to members of an organization. A policy provides direction to employees and is authoritative in nature. Standards document criteria such as a proven norm or method. They are typically external to an organization but can influence the organization's policies, guidelines, and procedures. Guidelines provide recommendations to members of an organization but aren't mandatory or authoritative in nature. Procedures provide individuals with specific action steps to accomplish tasks. Figure 9-5 shows the relationship among policies, standards, procedures, and guidelines.



EXAM TIP Policies are authoritative in nature and provide overall direction. Standards are established norms, which can influence policies. Guidelines and procedures are derived from policies. Guidelines provide recommendations but are not authoritative.

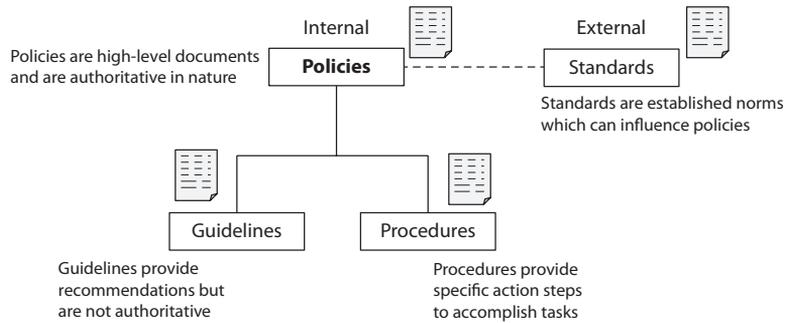


Figure 9-5 Relationship among policies, standards, procedures, and guidelines

The following are some examples of policies:

- **Security policy** Many organizations have security policies that define the overall goals of security. A security policy defines what security means to an organization and outlines basic security requirements. Personnel within an organization create standards and procedures to meet the goals and requirements specified within a security policy.



NOTE Chapter 12 covers security policies in more depth.

- **Acceptable use policy** This lets users know what they can and cannot do with IT systems owned and controlled by the organization. For example, some organizations allow users to send and receive personal e-mails, while other organizations restrict this.
- **Backup policy** Many organizations use a backup policy to define the data that should be backed up and how long backups should be retained. A backup policy would also include a statement mentioning that a copy of backups must be kept in an offsite location to protect against a potential disaster.

An organization may use configuration control procedures to dictate the configuration of systems. For example, the organization may identify specific baseline configurations for servers and baseline configurations for desktop systems. At any time, the organization can compare the current configuration of these systems against the baseline to ensure that the system is still meeting the requirements of the standard.

An organization can often choose what standards to follow, so standards aren't necessarily authoritative in nature. For example, the International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) 20000 (previously

known as BS 15000) is a standard based on the Information Technology Infrastructure Library (ITIL). An organization can use this as an external standard when creating its security policy. Management would ensure that the security policy addresses relevant elements of ISO/IEC 20000. Similarly, organizations can use NIST special publications in the SP 800 series as standards when creating or updating their security policies.

Guidelines are more generic in nature. For example, a guideline related to backups may recommend backing up all databases. Notice this is different from having an authoritative backup policy that states that all customer databases *must* be backed up.

Policies are implemented with procedures. For example, written procedures identify how users acknowledge the acceptable use policy. It may specify that employees need to review and sign the policy when hired and annually thereafter. Backup procedures would itemize the steps required to back up different types of data, such as a database or individual files, and the steps required to restore backups.

Response Plans

Many organizations prepare for disasters by creating different types of plans. A business continuity plan (BCP) helps an organization prepare for emergencies that can interrupt the mission of the business. The BCP uses a business impact analysis (BIA) to identify critical functions for an organization and then includes plans to keep these critical functions operating. A disaster recovery plan (DRP) identifies steps to recover critical systems after a disaster.

Security incidents can cause a significant amount of damage, so organizations often create incident response plans to prepare for them. An incident response plan provides detailed guidance for personnel responding to security incidents.



NOTE Chapter 12 covers disaster preparation in more depth, including the differences between a BIA, BCP, and DRP. Chapter 13 covers incidents and incident response in more depth.

Change Control and Configuration Management

Change control helps prevent unintended outages from occurring as a result of changes. There are countless stories of well-meaning technicians and administrators modifying systems to resolve one problem, but inadvertently creating another problem. In the worst-case scenarios, they end up disabling an entire system or service.

Instead of making a change on impulse when a change is required or desired, administrators submit a change request. A group of experts on a change review board reviews the request and tries to predict potential problems from the change. They approve the change if they predict it will not cause any problems, and reject the change if they determine it will be detrimental to operations.

Configuration management ensures that information about a system's configuration is available and helps ensure that similar systems are configured similarly. Many organizations deploy systems using baseline configurations with images. In addition to ensuring

that the system is deployed securely, configuration management ensures the system can later be checked and compared to the baseline to ensure that it hasn't been reconfigured.

Administrators develop and maintain system and security control documentation as a part of their daily tasks. The documentation can be contained within a change control system, a configuration management system, or a combination of both. If systems are compromised in an attack or from a disaster, this documentation is invaluable when returning the system to the state it was in prior to the incident.

Some organizations combine change control and configuration management into the same process or control, while larger organizations sometimes separate it into separate processes.



NOTE Chapter 10 covers change management and configuration control in more depth, including the use of imaging to implement baselines.

Testing Patches, Fixes, and Updates

As mentioned previously, vendors regularly release patches, fixes, and updates. Basic hardening practices ensure that systems have all relevant patches installed, including patches for operating systems and applications.

Some small organizations configure systems to apply all patches as soon as the vendor releases them. However, midsize and large organizations implement a patch management process. When a vendor releases a patch, administrators in these larger organizations often take the following steps, described in the following sections:

1. Evaluate the patch to determine if it is needed.
2. Test the patch to verify that it doesn't cause problems.
3. Apply the patch to systems that need it.
4. Audit systems to verify that the patch is applied.
5. Document the patch in a change management system.

Evaluating Patches

It's important to note that it isn't necessary to apply all patches to all systems. For example, a patch for a Microsoft Server product that implements a fix for Internet Information Services (IIS) is relevant for a server running the IIS web server application. However, it is not necessary or desired for a system not running IIS. With this in mind, administrators first evaluate patches, fixes, and updates to determine what systems need them.

Testing Patches

Unfortunately, some fixes cause problems elsewhere. In the worst-case scenario, systems no longer boot after administrators apply a patch. If this happens to one or two desktop computers, it can be inconvenient, but if it happens to a critical server or hundreds of desktop computers, it can be catastrophic. Because of this, organizations test the patches before applying them to critical systems or to large numbers of computers.

Ideally, the patch won't cause any problems. However, many patches have caused problems in the past. In one case, a hardware vendor used a specially configured image on a huge quantity of computers that it sold. Using images is common, but due to how the vendor created this image, an update caused these computers to go into an endless reboot cycle. Organizations that owned computers with this image had to recover each system individually after applying the patch. However, organizations that tested the patch on a single computer discovered the problem and did not apply the patch to other computers.

Some problems created by patches aren't as extreme. For example, one patch prevented audio from working on a specific brand of computers. A smaller organization that I know of configured its computers to receive patches automatically (without testing). One of the patches disabled all audio on about 50 of its 80 computers. However, not everyone complained right away. Instead, the IT team (one full-time IT professional and one part-time intern) started getting random complaints. It took the team a couple of weeks to realize how many computers the patch affected and identify the source of the problem.

Some organizations use IT administrator computers for testing. In other words, after determining that systems need a patch, they deploy it to computers used by IT administrators. If the patch causes any major problems, administrators are able to troubleshoot and recover from the problem more easily than a typical end user could.

Applying Patches

If administrators determine that the patch is needed and doesn't cause any unacceptable problems, administrators apply the patch. Many organizations use server applications to automate the deployment of the patches. Microsoft provides Windows Server Update Services (WSUS) as a free download to manage Microsoft updates. System Center Configuration Manager (SCCM or ConfigMgr) is a server application that organizations can purchase from Microsoft. ConfigMgr includes the features of WSUS plus many more, such as the ability inventory systems and push out applications. Additionally, third-party applications are available that organizations can use to manage updates.

After applying the patch, these automated tools verify that the systems have accepted it. In some cases, a system will block the installation of the patch. For example, a firewall can block patches if it isn't configured with an exception to accept them. It's also possible for an existing application to cause a conflict and block the patch.

Auditing Systems

Automated deployment applications such as WSUS provide administrators with a listing of all computers that have accepted the patch and a listing of computers that remain unpatched. Additionally, they periodically audit the systems to ensure that they all have approved patches. If a system is rebuilt or a patch is removed, these applications discover the problem. Administrators can configure WSUS and similar third-party applications to redeploy the update when they discover an unpatched system.

Unpatched systems are often vulnerable to attacks, so many vulnerability scanners check systems for patches. A vulnerability scanner can't deploy a patch, but it can provide an alert to security administrators that a patch is missing.

Documenting Patches

As mentioned previously, most organizations use a change management process to document any type of changes. This includes changes from patches, fixes, and updates. When administrators determine that a system needs a patch, they follow the change management procedures to apply the patch.

The change management system helps prevent unintended outages due to unapproved changes. It also provides documentation on approved changes. This documentation can be extremely useful if administrators need to rebuild a system from scratch.



EXAM TIP Whenever possible, it's best to test patches before applying them. As a follow-up step, administrators should audit systems to verify that patches are applied. Additionally, personnel should document applied patches in a change management system.

Endpoint Device Security

Many organizations implement controls to protect against risks associated with endpoint devices. In this context, endpoint devices refer to devices used by end users. They include mobile devices (such as smartphones and tablets), USB devices, laptops, and desktop computers.

Chapter 3 includes a section on protecting mobile devices, such as smartphones and tablets, and some of the risks that can occur if they are lost. Common controls with these devices are password protection, encryption of data, remote wipe, and GPS capabilities. If devices are lost or stolen, these controls help reduce risks.

Other times, the risks are due to built-in cameras or phones transmitting audio without the user's knowledge. Because of these risks, many organizations prohibit the use in certain areas of many small devices, such as cell phones. Some organizations have mini-lockers close to the entrances to secure areas and require employees to lock their mobile device in the locker before entering, and retrieve it when they leave.

Bringing Your Own Device

A trend in recent years is that employees want to bring their personal smartphones and tablets to work and connect into the work network. Bring your own device (BYOD) provides great benefits for organizations, but also presents some challenges.

Some of the benefits include reduced costs for organizations, because they don't need to provide these devices, and improved productivity for workers, because they remain connected to the work network even when they aren't at work. For example, employees can configure their devices with their work e-mail and be available to answer e-mails during nonwork hours. Of course, this availability brings with it the challenge of maintaining a work-life balance.

Keeping these devices up to date is another challenge. Organizations have processes in place to keep organization-owned devices up to date with current patches and updates. However, pushing updates to employee-owned devices presents other problems. First,

because updates might change how the device operates or how applications function, employees might not want to update their devices to a new operating system or with a new patch. Second, there are many different types of mobile operating systems (such as Apple's iOS and the Android operating system).

Android is the most popular operating system used worldwide on smartphones and tablets. However, it isn't run on a single brand of device. The Android operating system is open source, and multiple companies use it on their devices. For example, Google uses Android on many mobile devices that it sells. Samsung is one of many other companies that uses the Android operating system on multiple smartphones and tablets. However, because Android is open source, many mobile device vendors modify it for their devices. In order for an IT department to keep BYOD devices up to date, they need to be able to manage all the different versions of Android. This includes all the Google devices, Samsung devices, and the different devices manufactured by other companies using Android. This requires knowing which devices users have and then keeping up with the updates for all of them.

Another challenge is managing the data. The user owns the device. Who owns the data? Also, how is the organization's data protected when it is on the user's device? Organizations often implement BYOD policies to address these challenges and include elements of the BYOD policies in acceptable use policies.

Using Corporately Owned Devices

An alternative to BYOD is corporate-owned, personally enabled (COPE) devices. These are smartphones and tablets that the organization purchases and issues to employees for work purposes. A huge benefit of COPE devices is that an organization can control what devices it purchases. For example, instead of managing hundreds of different devices owned by employees, an organization can purchase only one or two models and issue these to employees.

Telecommuting

Many organizations support telecommuting (or telework), allowing employees to work from home. As long as employees have the discipline to be productive, telework is beneficial for both the employees and the organization. However, there are some risks.

Teleworkers typically access the organization's network using a remote access solution. Chapter 4 covers remote access solutions in more depth. A key element of remote access security is to use strong tunneling protocols such as Internet Protocol security (IPsec).

Another concern are the employees' computers. Because employees typically own the computers in their homes, the organization doesn't control them. However, if malware has infected an employee's home system, it is possible for an attacker to gain access to the organization's internal network via the employee's computer. Organizations mitigate this risk by using network access control (NAC) methods (discussed in Chapter 4). NAC methods inspect computers for predefined health characteristics, such as up-to-date AV software and enabled firewalls. Computers that don't meet these health requirements are isolated from the main network until they meet them.

Mobile Device Management

Organizations typically have strong patch management processes in place, as discussed earlier in this chapter. Similarly, they have automated methods to configure systems, push out applications to systems, and update these applications when necessary. Historically, technical methods only accessed internal computers owned by the organization. However, with the explosion of BYOD and COPE devices, organizations now have a need to implement similar methods for smartphones and tablets.

Mobile device management (MDM) solutions provide these services. As an example, Microsoft's ConfigMgr (mentioned in the "Applying Patches" section earlier in this chapter) includes management services for some mobile operating systems, including iOS and Android. Administrators use ConfigMgr to deploy updates, applications, and configuration settings.

It's worth repeating that hardware vendors, such as Samsung and Motorola, typically modify the open source Android operating system to meet their own needs. In other words, ConfigMgr (or any MDM system) cannot necessarily modify any device running Android, but instead needs software specifically designed for the BYOD or COPE devices.



TIP You can visit this site for a list of Android devices: www.androidtapp.com/list-of-android-devices/. Scroll down to the "Phones" or "Tablets" sections and you'll see that devices are running many different operating system versions. Additionally, these vendors can update the operating system at any time. This becomes extremely challenging for administrators to manage without MDM software.

USB Devices

USB devices, such as USB flash drives and USB hard drives, are especially troublesome to many organizations. It's easy to copy a significant amount of data to these devices, and it's also easy to lose them. Lost USB devices can result in a significant amount of data loss or data leakage. Some organizations require the encryption of all data stored on any type of device, including USB devices, to protect against the loss of confidentiality.

Of course, another risk related to USB devices is the risk of transferring malware from system to system. If a USB device is infected with a virus, that virus can easily infect a system as soon as a user plugs the USB device into a computer. Similarly, if a system is infected with a virus, the virus can detect when a user plugs the USB device into the computer and infect the USB device.



EXAM TIP Two significant risks related to USB devices are data leakage and infection with malware. Losses due to data leakage can be mitigated by encrypting data at rest. Malware infections can be mitigated with antivirus software.

Thin Clients

The price of computers is low enough that it's common for each user to have his or her own computer. However, in some cases, it's better to use a thin client to reduce some

risks. A thin client has very little hardware and only minimal software installed on it. It connects over a network to a server for all its data and applications. If the thin client is stolen, no data is lost. In contrast, a computer with an operating system and applications is a thick client.

Virtualization

Virtualization was mentioned in Chapter 4 with cloud computing. It is becoming more and more common today, and not just with cloud computing. Many organizations are realizing that they can get some significant benefits from virtualizing servers.

For example, an organization may have 500 servers hosted in a data center. The organization could instead purchase 50 very powerful servers, with each server hosting 10 virtual servers. Although the host servers are more expensive initially, they cost less in the long term. It's cheaper to provide power and cooling to 50 servers than to 500. Virtualizing a data center also requires less space and less physical security requirements.

One of the biggest risks with virtual machines (VMs) comes from VM escape attacks. VM escape attacks give an attacker access to the physical system from within a VM. In other words, the attacker “escapes” from the VM. When successful, the attacker can gain full control over all the VMs on the host in a successful VM escape attack. The primary protection against VM escape attacks is to keep the host system and all VMs up to date with current patches.

Application Whitelisting and Blacklisting

Application whitelisting includes a list of approved applications. When implemented, it prevents users from installing or running any applications that are not on the list. Blacklisting includes a list of nonauthorized applications. It prevents users from installing or running any application on the list. Some MDM applications support both whitelisting and blacklisting. Administrators create a list and configure the MDM application to use it. The MDM application can then scan devices to ensure they are in compliance.

Apple maintains the App Store, which is effectively a massive whitelist. Apple employees review all apps before publishing them in the App Store. Unless users jailbreak their Apple devices to bypass security, they can only install apps from the App Store. An exception to this is that an organization can create apps for its employees and configure Apple devices to accept these apps.

Endpoint Encryption

Many devices support whole disk or full device encryption. This helps prevent attackers from accessing data on stolen or lost devices. As an example, the iOS operating system encrypts all data on the device. This includes text messages, photos, contacts, notes, and more. The data is encrypted automatically as long as the device includes a screen lock, which the iOS enables automatically. Similarly, Android devices support encryption, and devices with Android 5.0 (Lollipop) or newer automatically encrypt user data.

Jailbreaking and Malware

Jailbreaking bypasses security and grants the user root-level access to the iOS operating system. It is similar to rooting a device running an Android operating system. It voids the warranty, so it isn't commonly done in most countries. Interestingly, it is common in China.

Free jailbreak tools also install apps for Chinese app stores such as Taig and Cydia. These app stores host apps similar to Apple's App Store. However, the owners of these alternative app stores receive the revenue instead of Apple.

Alternative (non-Apple) sites have hosted pirated software and software that included malware so they do not provide the same reliable whitelisting services that you'll find on Apple's App Store. As an example, WireLurker is malware that targets jailbroken Apple devices. It collects sensitive information (such as user's call logs, phone book contacts, and more) from the user's device. Users have downloaded apps infected with WireLurker more than 350,000 times from non-Apple sites such as Taig and Cydia.

WireLurker also infected some non-jailbroken devices via infected Apple Mac computers. The Macs were infected from apps downloaded from another alternative site, the Maiyadi app store. When users connect their iOS devices to the infected Mac using a USB cable, the malware attempts to install itself on the iOS device.



TIP Apple previously had a backdoor to encrypted data, which represented a potential risk. Apple personnel could unlock any device with a key they controlled. However, since the release of iOS 8, Apple no longer includes this backdoor key. This ensures iOS 8 isn't accessible to potential attackers too. Similarly, Android devices create the encryption key on the device, and this key isn't accessible by the manufacturer.

Trusted Platform Module

Many laptop computers support disk encryption using a Trusted Platform Module (TPM), which is a hardware chip embedded in the computer's motherboard. If the laptop includes the TPM, users can enable it via the operating system. However, if the laptop didn't ship with a TPM, you'd need to replace the motherboard to add one. The TPM stores cryptographic keys, which the operating system can use to encrypt and decrypt data.

The TPM includes a unique RSA asymmetric key (called an endorsement key), which the manufacturer embeds into the TPM before shipping it. When users enable the TPM, it creates a storage root key, which encrypts application keys. Applications (such as Microsoft BitLocker) create application keys to encrypt an entire disk. Note that it doesn't encrypt the entire disk drive. Instead, it partitions the disk drive into separate disks. The BitLocker partition remains unencrypted, but BitLocker encrypts the other partition where the data resides.

Sandboxing for Secure Browsing

Many web browsers (such as Google Chrome and Apple Safari) support sandboxing, which ensures that web pages and web applications have limited access to your system. Think of a sandbox in a children's playground. Parents can place their children in the sandbox and let them play. They know their children will remain safe as long as their children stay in the sandbox, which is isolated from any risks outside of the sandbox.

Similarly, web browsers run within isolated areas of a computer, which function as a virtual sandbox. The sandbox prevents web applications from accessing your personal files or hardware such as your webcam. Additionally, web browsers have limited permissions within the sandbox. These limited permissions are sometimes called protected mode or low integrity mode. Even if a user visits a malicious website that downloads malicious code, the malicious code isn't able to do any damage without first breaking out of the sandbox.

User Awareness and Training Programs

User awareness programs attempt to help users within the organization understand security procedures and practices. These programs help users understand their responsibilities related to security, recognize risky behaviors, and change their behavior if they are engaging in unsafe computing practices.

For example, without security training and education programs, many users don't understand the tactics used by social engineers and may readily give up important information to them. Similarly, users may not understand the risks of jailbreaking a device and installing apps from other app stores. Training topics are sprinkled through this book to stress the importance of awareness and training programs.

Understanding Fault Tolerance

A single point of failure is any component within a system that may cause a significant system outage if it fails. Fault-tolerant controls help eliminate outages from single points of failure by adding redundancies. If a system develops a fault, the system can tolerate the fault and continue to operate. Organizations often add fault tolerance for disks, servers, connections, and sites.



NOTE This section covers fault tolerance for disks, servers, and connections. Chapter 12 covers alternative sites and how they can provide fault tolerance for entire locations.

Fault Tolerance for Disks

A Redundant Array of Independent Disks (RAID) uses multiple disks to improve the performance of the disk subsystem, and most RAID arrays provide fault tolerance for the disks. You can implement RAID with some types of software or with hardware. For example, some operating systems have built-in support for software RAID. Some

motherboards have built-in support for hardware RAID, and you can purchase external hardware RAID disk subsystems.

Whenever possible, use hardware RAID to get better performance. In software RAID, the operating system has to provide the services to support the RAID and that can slow down overall system performance. Some of the common RAID arrays are RAID-0, RAID-1, RAID-5, and RAID-6.

RAID-0

RAID-0 uses two or more disks in an array to improve the performance of both read and write operations. However, RAID-0 does not provide any fault tolerance. If one drive in a RAID-0 disk subsystem fails, the entire array fails.

RAID-1

RAID-1 uses two disks in a mirror configuration, as shown in Figure 9-6. Each file is written completely to each drive in the array. If one of the drives ever fails, the other drive still includes complete copies of all of the files.

Most RAID-1 disk subsystems also include the ability to read from both drives at the same time. This improves read performance because the disk controller can read half the file from one drive while simultaneously reading the other half of the file from the other drive.

Hardware RAID-1 disk subsystems have the ability to recover automatically when a drive fails. For example, if one of the drives fails, the system recognizes the failure and stops using the failed drive. However, because all the data is on the drive that didn't fail, the system continues to operate without interruption. Of course, if one of the drives fails, the array is no longer providing fault tolerance. Administrators should replace the failed hard drive as soon as possible.

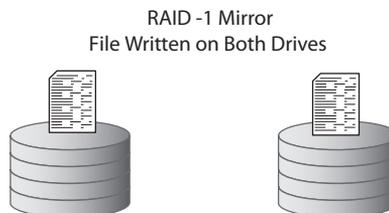


EXAM TIP A RAID-1 mirror provides fault tolerance by mirroring data on one drive to a second drive.

RAID-5

RAID-5 uses three or more drives in an array and uses striping with parity. It writes data onto the drives in stripes, and it uses the equivalent of one drive for parity data. If any single drive fails, the array can use the parity information to reconstruct the data and the array continues to operate.

Figure 9-6
RAID-1 mirror



Decimal Value	Data Bit (2 ¹)	Data Bit (2 ⁰)	Odd Parity Bit
0	0	0	1
1	0	1	0
2	1	0	0
3	1	1	1

Table 9-2 Odd Parity

RAID-5 uses either odd or even parity to give the number of bits in a stripe either an odd or even value. For example, if odd parity is used and the number of 1 bits in the stripe is even (such as two 1 bits or four 1 bits), then the parity bit will be a 1 to make the total number of 1 bits odd (such as three or five). If odd parity is used and the number of 1 bits in the stripe is odd (such as one 1 bit or three 1 bits), then the parity bit is a 0 to keep the total number of 1 bits odd.

Consider Table 9-2. It shows the decimal values of 0 through 3 and their binary equivalent represented as data bits. In the second data row (considered a stripe in RAID-5), a decimal value of one (1) has the data bits of 0 and 1, giving an odd number of 1 bits. The parity bit is 0, so the number of 1 bits remains odd. In the fourth stripe with a decimal value of three (3), the data bits are 1 and 1, which is an even number of 1 bits, so the parity bit is a 1.

Now imagine that each of these columns represents a disk drive and each of the rows represents a stripe of data, as shown in Table 9-3. The equivalent of two drives carries the data, while the equivalent of one drive carries the parity bits.

If any of these drives fails, the system can still calculate the contents of the drive simply by calculating the remaining data. If the remaining data in a stripe has an odd number of 1 bits, the missing data in the stripe is a 0. Similarly, if the remaining data in a stripe has an even number of 1 bits, the missing data in the stripe is a 1.

Table 9-4 shows the same RAID-5 array with a failed drive. Even if you don't know what the data was originally but you do know that the RAID is using odd parity, you can easily calculate the data in each of the stripes. This allows a RAID-5 array to suffer the loss of an entire drive and continue to operate.

	Data Bit (2 ¹)	Data Bit (2 ⁰)	Odd Parity Bit
Stripe	0	0	1
Stripe	0	1	0
Stripe	1	0	0
Stripe	1	1	1

Table 9-3 RAID-5 with Odd Parity

	Data Bit (2 ¹)	Failed Drive	Odd Parity Bit
Stripe	0	Unknown	1
Stripe	0	Unknown	0
Stripe	1	Unknown	0
Stripe	1	Unknown	1

Table 9-4 RAID-5 with Odd Parity and a Failed Drive

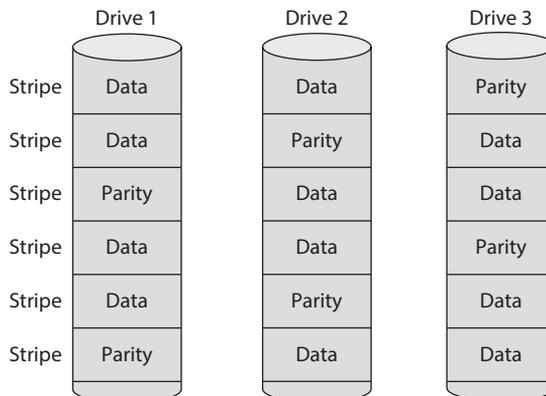


TIP RAID-5 stripes hold more than a single bit in each stripe, but for illustration purposes, it's much easier to show the values as single bits. RAID-5 stripes are actually 64KB.

It's important to stress that while a RAID-5 has the equivalent of one drive devoted to parity, the parity data is not written on a single drive. Figure 9-7 shows a three-drive RAID-5 array with multiple stripes. Notice that each stripe has two drives holding the data, with one drive holding the parity, but the parity data for different stripes is contained on different drives. Early implementations of striping with parity put the parity bit on a single drive, but engineers discovered that the drives performed better when the engineers spread the parity bit across multiple drives.

Both software-based and hardware-based RAID-5 disk subsystems can automatically detect when a drive fails and automatically continue to operate. However, hardware-based RAID-5 disk subsystems often have extra spare drives in the array. If a drive fails, the system detects the failure, logically removes the failed drive, logically adds a spare drive into the array, and rebuilds it. This all occurs automatically without any human intervention. The array also provides some type of notification to administrators of the failure so that they can replace the failed drive.

Figure 9-7
RAID-5 striping
with parity



RAID-6

RAID-6 is an extension of RAID-5. The formal definition of RAID-6 is a RAID that uses two parity blocks, though the parity blocks can be implemented differently. As a simple example, the second parity block can be a copy of the first parity block but stored on a different disk. The benefit is that a RAID-6 can survive the failure of two disk drives. RAID-6 requires a minimum of four disks, with the equivalent of two disks dedicated to parity.

Many storage vendors recommend the use of RAID-6 over RAID-5 today. This is partly due to the unrecoverable error rate (URE) of standard SATA drives and the size of current drives. As an example, think of a seven-disk RAID-5 with each disk 2TB in size. While UREs are relatively low, one URE can prevent the reconstruction of a RAID-5 disk subsystem. Further, the chance of a single URE from the 12TB of storage on the remaining six drives after a hard drive failure is as high as 62 percent.

However, with RAID-6, the disk subsystem will survive the failure of a single hard drive and a URE. As long as administrators quickly replace the single failed hard drive, they will be able to re-create the RAID.

Interestingly, most vendors gloss over the URE problem and instead stress that the RAID-6 can survive the failure of two hard disks. Technically, it can continue to operate with two failed disks. However, in reality, the disk subsystem is likely to experience a URE with two failed hard disks, and administrators will not be able to rebuild the array.



TIP More advanced RAID implementations are also available. For example, RAID-10 combines the benefits of a RAID-1 mirror with a RAID-0 striped array and uses at least four drives. This provides both increased performance and increased fault tolerance.

Failover Clusters

A failover cluster provides fault tolerance for one or more servers. A simple two-node failover cluster ensures a service continues to operate even if a server fails. For example, if an organization is hosting a database and wants to ensure that the database is always available, even if a server fails, the organization could use a failover cluster.



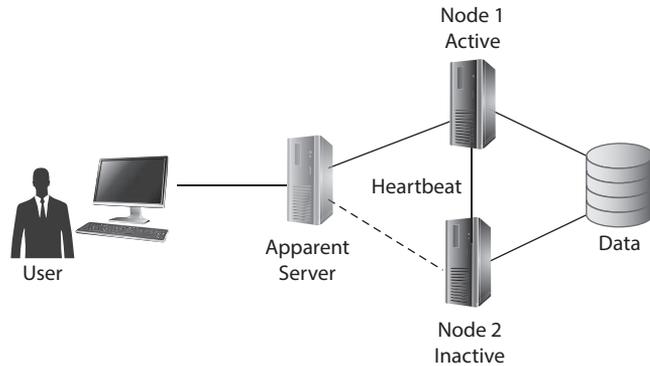
EXAM TIP Failover clusters allow a service to continue to operate even if a server fails. A failover cluster provides fault tolerance at the server level.

Figure 9-8 shows a simple two-node cluster with Node 1 active and Node 2 inactive. In this configuration, Node 1 is running the database management software and accesses the database on an external drive. When users query the database, they contact an apparent server for the cluster, and the clustering software routes the traffic to the active node.



NOTE Node 1 and Node 2 are both physical servers that you can touch. However, the apparent server is only an outward representation of the failover cluster. The clustering software presents this outward representation to users and internally routes all traffic to the active node.

Figure 9-8
Two-node
failover cluster



Node 2 in this cluster is currently inactive. It is running the same database management software and has access to the same database on the external drive. More importantly, it is constantly monitoring the health of Node 1 using a heartbeat signal. If Node 1 ever fails, Node 2 will know about it immediately. Depending on how administrators configure the failover cluster, Node 2 may try to restart services on Node 1 or take over services as the active node for the failover cluster.

If Node 2 takes over the services, the cluster software reroutes all requests to the failover cluster to Node 2. At this point, Node 1 is inactive and Node 2 is active. Users may notice a momentary glitch, but overall there is no loss of data.

Although a two-node cluster is easier to illustrate and conceptualize, clusters can be much more complex. For example, a two-node cluster can support two services instead of just one, with each node being active for one service and inactive for the other service. If either of the services fails, the other node can take over for both services. As another example, an eight-node cluster can support multiple services, with each node hosting one or more services and every other node capable of supporting the additional load if any other server fails.

Load-balancing clusters are another type of cluster, but they do not provide fault tolerance. Instead, a load-balancing cluster allows you to add more and more servers to balance a larger load. For example, an e-commerce site may be extremely busy during the holiday season. A load-balancing cluster allows the organization to add more servers to handle the additional website load. When the holiday season is over, they can remove the additional servers.

Redundant Connections

If an organization must stay connected to the Internet, it can implement two or more connections to ensure connectivity. Similarly, an organization can implement redundant connections to ensure that a main location remains connected to a remote office.

Organizations that use two or more connections often use one high-speed connection for primary use and a second, lower-speed connection as a redundant connection in case the first connection fails. For example, one connection might be a broadband connection,

and the second connection might be a slower fractional T1 connection (which is only a fraction of a full 1.54-Mbps T1 connection).

Understanding Backups

The importance of backing up data can't be overstated. Data will be lost at some point as long as an organization is using computers. Hardware failures, software bugs, accidental deletions, natural disasters, and malicious attacks can all destroy or corrupt data. However, the difference between a minor inconvenience and a major catastrophe is the existence of a backup. If a reliable backup exists, personnel can restore the lost data. If a backup isn't available, the lost data can result in significant losses.



EXAM TIP Backups are useful only if they can be restored. Organizations typically include test restore procedures in their backup plan. A test restore goes through the steps to restore data from a backup and verifies the usefulness of the backup.

A backup policy identifies what data is important to an organization and what data personnel need to back up. Backup policies include retention policies that identify how long the organization should keep backed-up data. These policies are authoritative in nature and can help justify the purchase of backup materials. Tape drives, tapes, backup software, and dedicated backup servers can all add up to significant costs. However, if a backup policy dictates the need for the backups, the policy provides justification for the costs.

In contrast, without a backup policy, administrators will make their own decisions on what to back up based on their own experience. These decisions are often driven by recent data losses. For example, if a high-ranking user recently lost data, then administrators might decide that backing up user data is most important. However, if no user data is permanently lost over a year's timeframe, administrators might decide that backing up user data isn't as important. Also, without a backup policy, administrators may never test these backups until someone loses data. If the backups aren't reliable, this is too late. The data is lost. Worse, when administrators ask for funds to replace worn-out tapes, management might decide not to approve the purchase.

A backup policy also stresses the importance of keeping a copy of backups in a separate geographical location. If backup tapes are kept with the server and a fire in the server room destroys the server, it will also destroy the backups. Even if the backups are stored elsewhere in the same building, a tornado or flood or other disaster can destroy all the backups. However, if the backups are stored offsite, they will be available even if a catastrophe destroys the site.

Traditional backups use full backups combined with either differential or incremental backups. These backups often store the backups on tapes, but the use of disks for backups is becoming more common. Backing up to tape is much slower, but it provides more data storage at an overall lower cost. As an alternative, many organizations use a disk-to-disk-to-tape strategy. This is useful when backing up online databases because the backup can slow down database performance. Administrators back up the data to a

disk, taking advantage of the fast disk speed. They then back up the data from the second disk to tape.

Full Backups

A full backup simply backs up the entire contents of the target data. This is useful if there isn't much data to back up or the data doesn't change very often. If the backups are small, you can back up the data every day.

However, most backups are rather large and it takes a long time to back up the data and a long time to restore it. Alternatives to full backups include full/incremental backup strategies and full/differential backup strategies. If you want to minimize the amount of time that it takes to back up the data, you can use a full/incremental backup strategy. If you want to minimize the amount of time that it takes to restore the data, you can use a full/differential backup strategy. Both of these strategies start with a full backup.

Full/Incremental Backup Strategy

In a full/incremental backup strategy, you back up all of the data to start the backup plan. Then on a regular basis you back up only the data that has changed since the last backup. For example, consider Table 9-5 as you read the following list, which describes the backups performed on each day:

- Full backups are completed early on Sunday morning at 1 A.M. using an automated system.
- At 1 A.M. on Monday morning, the incremental backup backs up only the changes made since the last backup (the full backup on Sunday, in this case).
- At 1 A.M. on Tuesday morning, the incremental backup backs up only the changes since the last backup (the incremental backup on Monday, in this case).
- At 1 A.M. on Wednesday morning, the incremental backup backs up only the changes since the last backup (the incremental backup on Monday, in this case).
- Throughout the week, the incremental backups would be about the same size.



NOTE Table 9-5 shows backups until Wednesday for brevity, but administrators would continue with the backups through the rest of the week (on Thursday, Friday, and Saturday). On Sunday the process starts again with a full backup.

Sunday Morning	Monday Morning	Tuesday Morning	Wednesday Morning	...
Full backup	Incremental Backs up changes since full on Sunday	Incremental Backs up changes since incremental on Monday	Incremental Backs up changes since incremental on Tuesday	...

Table 9-5 Full/Incremental Backup Strategy

If a data loss occurs, you can use the backups from Sunday and then use each of the backups from individual days to restore the data. For example, imagine that a hard drive failure occurs on Wednesday afternoon. You would restore the full backup from Sunday, then the incremental backup from Monday morning, then the incremental backup from Tuesday morning, and then the incremental backup from Wednesday morning.

Some of the recovered data would only be included in the full backup, while some of the recovered data would only be included in the backup from Tuesday, and some would only be included in the backup from Wednesday. Here a few examples:

- One person might create a report on a monthly basis, but did not modify it on Monday, Tuesday, or Wednesday. It would only be in the full backup.
- Another person might create a weekly report every Monday but not modify it on Tuesday, Wednesday, or any other day of the week. It would only be included in the Tuesday incremental backup.
- Help desk personnel might modify files seven days a week. Changes would be included in each incremental backup file (Monday morning, Tuesday morning, and Wednesday morning).



TIP A full/incremental backup strategy minimizes the time needed for backups during the week. However, the recovery can take longer because multiple backups typically need to be restored when recovering the data.

Full/Differential Backup Strategy

In a full/differential backup strategy, you back up all of the data to start the plan, and then on a regular basis you back up all of the data that has changed since the last full backup. For example, consider Table 9-6 as you read the following list, which describes the backups performed on each day:

- Full backups are completed early on Sunday morning at 1 A.M. using an automated system.
- At 1 A.M. on Monday morning, the differential backup backs up all the changes since the last full backup on Sunday.
- At 1 A.M. on Tuesday morning, the differential backup backs up all the changes since the last full backup on Sunday, including the changes from Monday.
- At 1 A.M. on Wednesday morning, the differential backup backs up all the changes since the last full backup on Sunday, including the changes from Monday and Tuesday.
- Throughout the week, the differential backups would steadily increase in size.

If a data loss occurs, you can use the backups from Sunday and then the most recent differential backup. For example, imagine data was lost due to a hard drive failure on

Sunday Morning	Monday Morning	Tuesday Morning	Wednesday Morning	...
Full backup	Differential	Differential	Differential	...
	Backs up changes since full on Sunday	Backs up changes since full on Sunday	Backs up changes since full on Sunday	

Table 9-6 Full/Differential Backup Strategy

Wednesday afternoon. You would need to restore the full backup from Sunday and then restore the differential backup from Wednesday morning.



TIP A full/differential backup strategy takes longer to back up during the week than a full/incremental strategy. However, the recovery time is reduced because a maximum of two backups are needed to recover the data.

Chapter Review

Controls include both countermeasures and safeguards, and security professionals use these interchangeably. They are the means, methods, actions, techniques, processes, procedures, or devices that reduce the vulnerability of a system or the possibility of a threat exploiting a vulnerability in a system. The primary goals of controls are to prevent, detect, and correct loss of confidentiality; loss of integrity; and loss of availability of a system.

Preventive controls attempt to prevent losses before they occur, detective controls attempt to detect violations either as they are occurring or after they've occurred, and corrective controls attempt to reverse the impact of the event. Compensating controls provide an alternative if circumstances prevent an organization from implementing a primary control. Deterrent controls encourage people to decide not to take certain undesirable actions that could result in an incident.

Control classifications include management (or administrative), technical, operational, and physical. Management controls focus on the management of risk and the management of IT security. Technical controls are implemented using technological methods, such as hardware, software, or firmware components. Operational controls focus on the day-to-day tasks that employees (rather than systems) perform. Physical controls refer to the controls you can touch.

System hardening is the practice of making a system more secure from the default settings and includes removing or disabling unused services and protocols, changing defaults, keeping systems up to date, enabling firewalls, and using AV software.

Policies are high-level documents used to provide overall guidance to an organization. Standards document criteria such as a proven norm or method and are typically external to an organization, but they can influence an organization's security policies. Guidelines and procedures are derived directly from the organization's security policy.

Business continuity plans and disaster recovery plans are examples of response plans that help an organization respond to a disaster. Change control procedures help prevent outages due to unauthorized changes. Configuration management procedures help ensure that systems are configured similarly using baselines. Patch management procedures evaluate patches, test relevant patches, and apply relevant patches that don't cause

unacceptable problems. Automated tools audit systems to ensure patches are applied, and change management systems document approved patches.

Endpoint device security typically focuses on mobile devices such as smartphones and tablets. If an organization allows employees to connect their personally owned devices to the organization's network, it will typically create a bring your own device (BYOD) policy and add elements of the policy to an acceptable use policy. Alternatively, an organization may opt for corporate-owned, personally enabled (COPE) devices to maintain more control of mobile devices. Mobile device management helps ensure that both BYOD and COPE devices are kept up to date.

User awareness and training programs raise the overall security posture of an organization by helping users understand their responsibilities associated with security and reduce risky behaviors.

Fault tolerance helps eliminate outages from single points of failure by adding redundancies. Redundant disk systems (such as RAID-1, RAID-5, RAID-6, and RAID-10) allow a system to continue to function even if a disk fails. Failover cluster configurations allow a server to continue to operate even if a server fails. Redundant connections allow an organization to stay connected to the Internet or to other locations, even if a connection fails.

Backups ensure that data is not lost even if the original data is lost due to a hardware failure, software bug, accidental deletion, natural disaster, or malicious attack. Backup policies identify what data to back up, how long to retain backups, and the importance of storing backups in separate geographical locations.

Backup types include full backups, incremental backups, and differential backups. Each strategy includes a full backup, which backs up all the important data. A full/incremental backup strategy completes a full backup and then backs up only the data that has changed since the last full or incremental backup. Incremental backups are about the same size and take the least amount of time to back up. However, they can take longer to restore because multiple incremental backups may need to be included in the restoration.

A full/differential backup strategy starts with a full backup and then backs up all the changes since the last full backup without regard to what was backed up during previous differential backups. Differential backups get steadily larger because they are backing up more and more data, requiring more time to complete as the week progresses. A full/differential backup strategy is the quickest to restore because only the full and the last differential backups are needed.

Questions

1. Which of the following provides the best definition of a control?
 - A. The means, methods, actions, techniques, processes, procedures, or devices used to prevent attackers from launching attacks on systems
 - B. A detective method that identifies threats
 - C. A corrective method that reverses the impact of an incident
 - D. The means, methods, actions, techniques, processes, procedures, or devices used to reduce the vulnerability of a system or the possibility of a threat exploiting a vulnerability

2. What are the primary objectives of security controls?
 - A. Prevent, detect, and correct
 - B. Prevent, detect, and block
 - C. Detect, correct, and block
 - D. Detect, correct, and remove
3. Which of the following controls attempts to avoid security incidents?
 - A. Preventive
 - B. Compensating
 - C. Corrective
 - D. Detective
4. What type of control is an audit log?
 - A. Technical
 - B. Corrective
 - C. Detective
 - D. Preventive
5. Which of the following security controls can restore a failed or disabled control?
 - A. Preventive
 - B. Corrective
 - C. Detective
 - D. Deterrent
6. What type of control are procedures to back up and restore data?
 - A. Operational
 - B. Corrective
 - C. Detective
 - D. Deterrent
7. A security professional is reviewing existing security controls. What type of security control is this?
 - A. Management
 - B. Technical
 - C. Physical
 - D. Compensating

8. Which of the following choices best describes an operational control?
 - A. A control implemented by people (rather than systems)
 - B. A control implemented using hardware, software, or firmware
 - C. A control that focuses on the management of risk and the management of IT security
 - D. A control that focuses on preventing losses due to risks
9. A computer system records events into a security log, and administrators periodically review the log for security incidents. What type of control is this?
 - A. A preventive, technical security control
 - B. A detective, physical security control
 - C. A preventive, physical security control
 - D. A detective, technical security control
10. Which of the following definitions best describes system hardening?
 - A. Making a system more secure than the default configuration
 - B. Increasing physical security to make it harder to access the system
 - C. Increasing the length of the administrator password to make it harder to access the system
 - D. Reducing the attack surface
11. Of the following choices, what is NOT one of the methods or goals of hardening a server?
 - A. Reducing the attack surface
 - B. Keeping a system up to date
 - C. Disabling firewalls
 - D. Adding AV software
12. Of the following choices, which provides high-level guidance to employees?
 - A. Procedure
 - B. Policy
 - C. Action steps
 - D. Disaster recovery plan
13. Of the following choices, which one is authoritative in nature?
 - A. Procedure
 - B. Policy
 - C. Action steps
 - D. Tasks

14. What is the overall goal of a change management process?
 - A. To slow down changes
 - B. To ensure that systems are configured similarly
 - C. To enable stakeholders to deny unwanted changes
 - D. To reduce unintended outages from unauthorized changes
15. What should administrators do after learning that a vendor has released a patch that is relevant for servers they manage?
 - A. Apply the patch.
 - B. Test the patch.
 - C. Audit systems to see whether the patch is applied.
 - D. Document the systems where the patch is applied.
16. Which of the following helps ensure that mobile devices have all relevant patches?
 - A. BYOD
 - B. COPE
 - C. MDM
 - D. USB
17. What is an important benefit to organizations that use virtual servers? (Choose all that apply.)
 - A. VM escape capabilities
 - B. Better control of data with cloud computing
 - C. Reduction of costs associated with power and cooling
 - D. Reduction of costs for physical security
18. You have two disk drives and you want to provide fault tolerance by mirroring the two drives. What should you use?
 - A. RAID-0
 - B. RAID-1
 - C. RAID-5
 - D. RAID-6
19. You need to ensure that a service continues to run even if a server fails. What should you implement?
 - A. RAID-1
 - B. RAID-6
 - C. Failover cluster
 - D. Warm site

20. You don't have enough maintenance time during the week to perform full backups, so you decide to implement a backup strategy that takes less time to do backups during the week. Of the following choices, what strategy will minimize the amount of time needed to restore a backup after a failure?
- A. Full
 - B. Incremental
 - C. Full/incremental
 - D. Full/differential

Answers

1. **D.** A control provides the means, methods, actions, techniques, processes, procedures, or devices that reduce the vulnerability of a system or the possibility of a threat exploiting a vulnerability in a system. You can't actually prevent attackers from launching an attack, but you can reduce their possibilities of success by either reducing vulnerabilities or reducing the impact of the threat. Controls can be preventive, detective, and/or corrective, but it isn't accurate to limit a control to only one of these types.
2. **A.** Controls attempt to prevent, detect, and/or correct losses to confidentiality, availability, or integrity. Some preventive controls attempt to block threats, and some corrective controls attempt to remove vulnerabilities, but neither blocking nor removing is referred to as a primary objective of security controls.
3. **A.** A preventive control attempts to avoid security incidents by preventing them. A compensating control provides an alternative when a primary control fails or is unavailable. A corrective control attempts to reverse the effects of an incident and can restore a failed or disabled control. A detective control attempts to detect incidents either as they are occurring or after they've occurred.
4. **C.** An audit log is a detective control because it identifies events either as they are occurring or after they've occurred. Technical is a class of control. A corrective control takes action to reverse the effects or impact of an incident. A preventive control prevents the event from occurring.
5. **B.** A corrective control can restore a failed or disabled control. A preventive control attempts to prevent incidents, and a detective control attempts to detect incidents. A deterrent control attempts to dissuade or deter personnel from trying to circumvent security policies or otherwise cause an incident.
6. **B.** Procedures to back up and restore data are corrective controls because they take action to reverse the effects of data loss. Operational is a class of control. Deterrent controls attempt to deter personnel from causing a security incident. A detective control identifies events either as they are occurring or after they've occurred.

7. **A.** Reviewing security controls is a management control. This is part of risk management, and reviewing existing security controls is part of a risk assessment. Technical controls use technical means, not an individual such as a security professional. Physical controls refer to the controls you can touch. Compensating controls are controls used as an alternative if the primary controls cannot be used. It's also worth mentioning that the other three answers all refer to classes of controls, while compensating controls refer to a control goal (similar to how preventive, detective, and corrective controls are control goals).
8. **A.** People, rather than systems, implement an operational control. A technical control is implemented with hardware, software, or firmware. A management control focuses on the management of risk and the management of IT security. A preventive control focuses on preventing losses due to risks.
9. **D.** This describes an audit log, which is a detective, technical security control. It is detective because it identifies incidents after they've occurred. It is technical because it is implemented with technology. It doesn't prevent the events, and the log is a digital file, not a physical item that you can touch.
10. **A.** Hardening a system is the practice of making it more secure than the default configuration. The other answers provide steps that can increase security on a system, but they don't provide an overall description of system hardening.
11. **C.** Disabling firewalls is not a method used to harden a server. However, administrators would enable firewalls to harden a server. All of the other answers are valid methods or goals of hardening a server. Administrators reduce the attack surface by disabling or removing unneeded services and protocols.
12. **B.** Policies provide high-level guidance to employees. Procedures include action steps to accomplish tasks. A disaster recovery plan can include one or more procedures to use in response to a disaster.
13. **B.** A policy is authoritative in nature and provides high-level guidance to employees. Procedures include action steps to accomplish tasks.
14. **D.** The overall goal of a change management process is to reduce unintended outages from unauthorized changes. The intent isn't to slow down changes, but that is often a side effect. A change management process gives stakeholders the ability to propose changes, but several entities within the organization grant or deny changes. A configuration management process ensures that systems are configured similarly.
15. **B.** Administrators should test patches before applying them to detect potential problems. If the patch is applied as soon as it is released, it can result in unforeseen problems. Systems are audited after the patch is applied to ensure that the patch was successfully applied. Change management and configuration control procedures are used to document patches.

16. **C.** Mobile device management (MDM) solutions help ensure that mobile devices (such as smartphones and tablets) have all relevant patches. Bring your own device (BYOD) refers to employees bringing their own devices to work and connecting them to an organization's network. Corporate-owned, personally enabled (COPE) devices refer to mobile devices that an organization purchases and issues to employees. Universal Serial Bus (USB) cannot apply patches to mobile devices.
17. **C, D.** Organizations often use virtualization to reduce costs associated with power and cooling and associated with physical security. Virtualization requires fewer physical servers to power, cool, and physically secure. Virtual machine (VM) escape is one of the biggest risks associated with virtual servers. Virtual servers can be used with cloud computing, but cloud computing reduces control of data.
18. **B.** RAID-1 provides fault tolerance by mirroring two drives. RAID-0 does not provide fault tolerance. RAID-5 uses three or more drives. RAID-6 is an alternative to RAID-5 and uses four or more drives.
19. **C.** A failover cluster includes two or more servers (called nodes), and if one server fails, the other server is able to handle the load for that server. RAID-1 and RAID-6 provide fault tolerance for disk drives. A warm site is a type of alternative site used to provide fault tolerance for an entire location.
20. **D.** A full/differential backup strategy takes the least amount of time to restore because it will require only two backups to restore: the full backup and the last differential backup. The scenario says that there isn't enough time to do a full backup during the week, and you can't do an incremental backup by itself. A full/incremental would require you to restore the full backup and each of the incremental backups since the last full backup, which usually requires restoring more than just two backups.

This page intentionally left blank

Auditing

In this chapter, you will learn about

- The value of auditing to enforce accountability
- Using clipping levels to define acceptable thresholds
- Using audit trails to re-create events leading to security incidents
- Types of audit logs used to create audit trails
- Using security audits to verify compliance
- Using baselines for configuration control
- Preventing outages with change management

Understanding Auditing and Accountability

Chapter 1 introduced the AAAs of security: authentication, authorization, and accounting. When used together, they ensure that only authorized entities (such as users and applications) can access systems or data and that a record of such activities exists. Chapter 2 presented access controls, including the three primary authentication factors, and different access control models. Auditing provides the accounting component, which tracks and records individual actions.

There are two primary methods used to perform auditing. They are different processes, but both methods are part of an overall auditing strategy:

- **Auditing activity through logs** Audit logs record activity and can be inspected by auditors at any time to reconstruct events. Logs may be on servers, firewalls, routers, or any other type of device.
- **Auditing activity through an inspection process** Periodic security audits inspect an organization to determine whether the organization is following required policies and procedures. For example, a security policy may state that administrators must use 15-character passwords, and a password audit would check the administrators' passwords to determine whether they are at least 15 characters long.



NOTE The “Exploring Audit Logs” section later in this chapter covers audit logs, and the “Performing Security Audits” section covers the second type of auditing using an inspection process.

Holding Users Accountable with Audit Logs

As long as the organization uses strong authentication practices, administrators, security personnel, and auditors can use the audit logs for accountability. As an example, Figure 10-1 shows a security log entry from Windows 7. The entry shows that a user named Bob read a file named ProjectMilestones.rtf in the C:\ProjectData folder using the WordPad application.

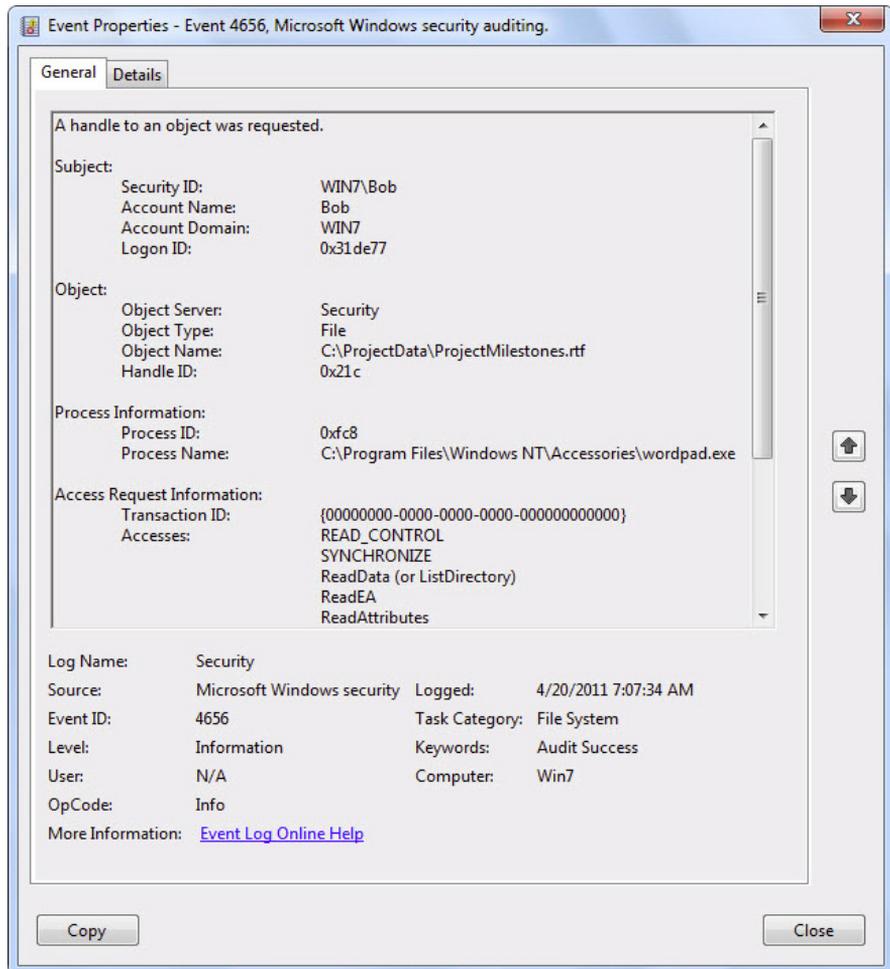


Figure 10-1 Windows 7 security log entry

It could be that Bob is a project manager and needs access to that file on a regular basis. It could also be that Bob is an administrator completely uninvolved in project management and abused his administrative privileges by accessing the file.

It's worth repeating that strong authentication is an important element here. If weak passwords are used, Bob could claim that he never looked at the file and suggest that someone guessed his password and impersonated him. On the other hand, if the organization implemented multifactor authentication and Bob was required to log on with a smart card and personal identification number (PIN), he won't be able to deny that he accessed the file. In this case, the audit logs provide nonrepudiation.

Other access control methods, such as video cameras and badge reader logs that record when someone enters or exits a secure area, provide added accountability. If these other methods documented that Bob was in the general location, they provide additional evidence for accountability.



EXAM TIP Audit logs combined with strong authentication and authorization practices provide nonrepudiation. When logs record specific user actions, users are unable to deny they took an action. At least, they don't have any credibility if they do try to deny taking the action.

Sometimes, a higher authority mandates specific accounting practices. For example, the Federal Information Processing Standards Publication 200 (FIPS Pub 200) published by the National Institute of Standards and Technology (NIST) specifies that federal agencies must meet at least the following minimum requirements:

- (i) create, protect, and retain information system audit records to the extent needed to enable the monitoring, analysis, investigation, and reporting of unlawful, unauthorized, or inappropriate information system activity; and*
- (ii) ensure that the actions of individual information system users can be uniquely traced to those users so they can be held accountable for their actions.*

Although this squeezes many words and phrases into a single long sentence, it does provide a good summary of the requirements of audit logs for federal agencies. It's also a good summary of the requirements any organization can use when implementing auditing and accountability.

Auditing with Logs

Audit logs record details such as who, what, when, and where for any events of interest. For example, Figure 10-1 (shown earlier) allows you to identify who took an action, what action the user took, when, and where:

- **Who** The local user account Bob on the computer named Win7 (Win7\Bob).
- **What** Bob read the ProjectMilestones.rtf file using the WordPad application.
- **When** The entry was logged on April 20 at 7:07 A.M.
- **Where** The file was located in the C:\ProjectData folder on the Win7 computer.

This example shows that Bob read the file. If Bob modified or deleted the file, the audit log could also show those actions, depending on how administrators configured logging.



TIP In Microsoft and most other computing systems, administrators must specify what to audit. If a system automatically audited all types of access to every single file, folder, and other system resource, the logs would become astronomically large. They would consume drive space and contain so much data that retrieving useful information would be difficult for administrators.

The primary purpose of auditing is to be able to reconstruct events leading up to and occurring during an incident. A secondary purpose is to serve as a deterrent. In other words, if users are aware that logs record their activities, they may be less likely to engage in unauthorized activities.

As an example, imagine an organization has an acceptable use policy (AUP) that states that users should only use computers and networking resources for work purposes. The AUP specifically prohibits accessing specific categories of websites, such as gambling sites. Non-work-related sites could easily be malicious and attempt to infect user systems with malware. However, if users know they shouldn't be visiting these sites and that logs are recording their activity, they may be less likely to visit these sites. The net effect is a reduction of incidents.

Proxy servers can block user access to unauthorized sites. Many organizations configure the proxy server to display a warning page when users attempt to visit unauthorized sites. These warnings often repeat the relevant information from the AUP. They also show the user information that the proxy server has logged, such as the username, the URL of the blocked site, and the date and time when the user attempted to access the site.



EXAM TIP Audit logging can reduce incidents simply by deterring users from engaging in unauthorized behaviors.

Clipping Levels

Many automated accounting systems generate alerts after the system detects a preset number of events. The system accepts or ignores these events until they reach a predetermined threshold, or *clipping level*. A classic example is a report based on failed logon attempts. Instead of reporting every single failed logon attempt, a system uses a clipping level to generate a report only if an account has exceeded a set number of failed logon attempts in a given period.

For example, administrators can configure an accounting system to generate a report if a log records three failed logon attempts in a 30-minute period for a single user account. If a user enters the wrong password once or twice, the system ignores those attempts as acceptable, normal user errors, and doesn't generate a report. However, if a user enters the wrong password three times within 30 minutes, it exceeds the clipping level, and the accounting system generates an alert.

Some automated auditing tools perform audit-reduction tasks. In other words, they have the ability to examine a log and retrieve relevant data from the log while ignoring the irrelevant data. These audit-reduction tools use clipping levels to identify events of interest based on predetermined thresholds.

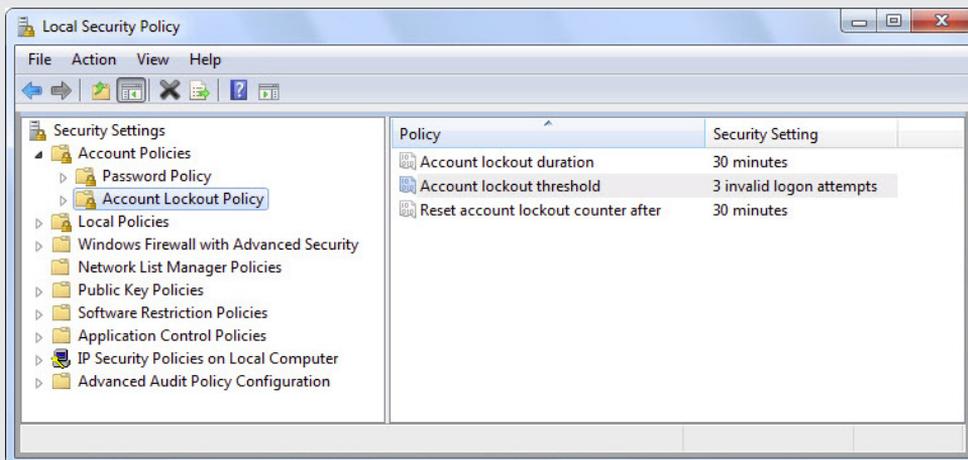


EXAM TIP A clipping level ignores events such as normal user errors until the auditing system reaches a predetermined threshold. However, the accounting system generates an alert after it detects the number of events has reached the threshold. Clipping levels such as this ignore individual logon failures and only take action after reaching the threshold.

Account Lockout

Many systems use an account lockout policy that locks out an account after a predetermined number of failed logons. These policies use clipping levels and only lock out the user after the number of failed logons exceeds a threshold.

For example, the following illustration shows an account lockout policy within the local security policy for a Windows system. It has the Account Lockout Threshold set to “3 invalid logon attempts,” so if a user enters the wrong password once or twice, the system doesn’t lock out the account. The Reset Account Lockout Counter After setting is set to 30 minutes, so as long as users don’t enter the wrong password three times within a 30-minute period, the system doesn’t lock the account. After 30 minutes, the system resets the incorrect password counter to zero.



You can also see that the Account Lockout Duration setting is 30 minutes. If the system locks out an account, it automatically unlocks the account after 30 minutes without any administrative intervention. If this setting is 0, the system locks the account until an administrator unlocks it.

Understanding Audit Trails

An audit trail includes one or more logs that track events occurring on a system or network. When auditors have access to all the logs, they are able to re-create the events that occurred leading up to the incident and determine what actually occurred during the incident. In other words, the audit logs allow auditors and administrators to re-create all the activity, or the trail of activity, created by an individual or process.

For example, if an attacker launched an attack on a web server in a demilitarized zone (DMZ), you could use firewall logs on the external firewall to identify the path the attacker used to get into the DMZ. You could then use the logs on the web server itself to track what the attacker did once he reached the web server. Logs on the web server could include host-based firewall logs, system logs, and security logs.

In some cases, the “who” identifies a specific account, such as a user account. In other situations, the “who” identifies the source of activity, such as an IP address. It’s worth noting that many attackers launch attacks through other locations or spoof the source IP address. In other words, the logged IP address doesn’t always provide the actual IP address of the attacker.

An audit trail is a technical detective control. However, when configuring or evaluating an audit trail, you should consider the following items to ensure it is effective as a detective control:

- **Does the audit trail provide a full accounting of user activity?** For example, does the audit trail record who, what, when, and where?
- **Is access to the logs controlled?** If anyone can access the logs, it makes it easy for attackers to erase the logs after an attack.
- **Are copies of logs stored on remote servers?** This ensures that even if an attacker erases or corrupts a log file after an attack, the remote server still holds a copy of the original logs.

Exploring Audit Logs

There are many different types of audit logs on computing and networking systems. Operating systems include regular logs that record activity on the computer; applications frequently record information in application logs based on the needs of the application; and network devices (such as routers and firewalls) regularly record different types of data to document network activity.

This section describes many of these logs. It also describes how administrators and auditors use these logs to track activity. Together, these logs can help an organization create a comprehensive audit trail documenting activity on individual systems and networks.

Operating System Logs

Operating systems include logs that track activity on individual computers and servers. Operating systems also provide the means to review these logs, such as the Event Viewer in Windows systems. Microsoft operating systems include at least the System, Application, and Security logs, but newer Windows operating systems include additional logs.

Security logs can record activity from any auditable activity and can record both success and failure events. In other words, if a user tries to open a file and successfully opens it, it is a success event. However, if the user doesn't have appropriate permissions, the system blocks the user's access and can record it as a failure event. Security logs don't record all activity by default, but instead require administrators to configure logging based on the organization's needs. For example, an administrator can configure object access logging for objects such as folders, files, or printers, and then when a user accesses one of these objects, the security log records the event. Other auditable security events include deleting or modifying files or folders, modifying system settings such as the system time, modifying permissions, and more.



TIP It's not useful to audit all object access. Instead, administrators choose to audit access to sensitive or proprietary data. For example, an organization may choose to record all access to classified research and development data, but not record access to publicly available data.

Some of the other logs used in Windows systems include

- **System log** Records events from the operating system, such as when the system boots or is shut down, and when a driver or service stops or starts.
- **Application log** Records events from specific applications. These events can be from end-user applications, such as a web browser, or server-based applications, which include database server applications and networking services, such as a server running the Domain Name System (DNS) service.
- **Setup log** Records events related to the setup of certain applications.

Microsoft has included additional features with logs in Event Viewer. For example, Figure 10-2 shows the Event Viewer in Windows 8 displaying some of the information from a successful logon event. Event Viewer allows you to view details from any of the captured events. You can also see Forwarded Events and Subscriptions in the left pane of this figure. Administrators create subscriptions on a central system that collects forwarded events from other systems. This allows administrators to monitor events from multiple systems in the network from a centralized computer.

If you manage 10 servers hosting Microsoft SQL Server databases, you can forward relevant events to an 11th system to monitor all of these servers. The original servers will still hold the original logs. However, the monitoring server hosts a copy of the logs in a remote location. Even if an attacker modifies the logs on one of the servers after an attack, the remote server still holds all the log entries within the Forwarded Events log.

Storing Logs on Remote Systems

Chapter 8 mentioned that attackers could modify logs stored on a local system (rather than a remote system). Administrators and investigators should view these logs with suspicion after an attack.

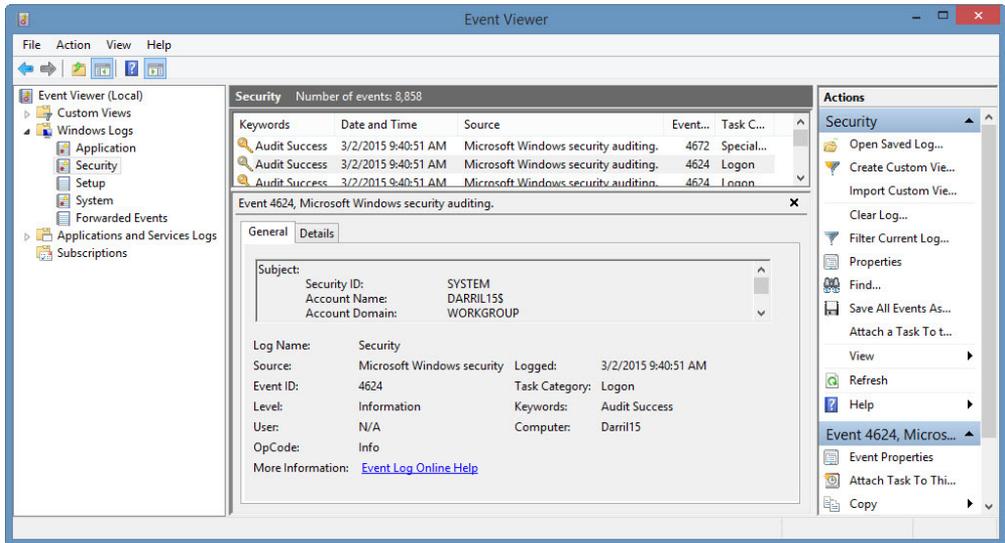


Figure 10-2 Microsoft Event Viewer

However, if you are able to store the logs on remote systems, it is much less likely that an attacker will be able to modify these remote logs. In other words, by storing logs on remote systems, you can help retain their integrity.



EXAM TIP Many attackers attempt to erase the audit trail by erasing or modifying individual logs. Their goal is often to get in, attack the system, erase or modify the logs, and then get out. However, logs on remote systems are more difficult for attackers to erase or modify.

*Nix Logs

UNIX and UNIX derivatives (such as Linux) are often referred to as *nix systems, and these systems have their own types of logs. The log file locations sometimes vary from one version to another, but there are some similarities. For example, system and accounting log files are usually in one of the following directories: `/var/log`, `/var/adm`, and `/usr/adm`.

Some of the common *nix log files and their contents include the following:

- **syslog** Many security processes and other utilities write to syslog. It's often a good place to view overall system activity.
- **sudo** This log records all attempts to use the su command. The su command (sometimes called the super user command) grants the user elevated privileges. Attackers often attempt to use su in various attacks, and uneducated users can misuse it.

- **auth** This log records authentication events on a Secure Shell (SSH) server. You can use it to determine whether a system intrusion has occurred. For example, a brute-force attempt to log in to the system will show multiple failed SSH login attempts, usually from multiple usernames.
- **maillog** The system records inbound and outbound mail activity in this log.

Proxy Server Logs

Most proxy servers can log all activity through the proxy server. This allows an organization to track websites that users access with their web browsers, how much time they spend on the sites, and all attempts to visit unauthorized sites. Chapter 4 presented proxy servers. As a reminder, an organization typically uses a proxy server for caching and URL filtering.

You may recall that an organization can use a proxy server to retrieve web pages from the Internet on behalf of internal clients. The proxy server uses Network Address Translation (NAT) to translate private addresses from internal clients to public addresses used by Internet-based websites. The proxy server caches web pages. If a user requests a cached web page, the proxy server sends the page to the user without retrieving it from the web server again. This reduces Internet bandwidth usage.

URL filtering prevents users from accessing unauthorized websites. If a user attempts to go to an unauthorized website, the proxy server blocks the request. As mentioned earlier in this chapter, proxy servers can also log these unauthorized attempts. When a user visits an unauthorized site, some organizations display a warning similar to that shown in Figure 10-3. When users see that the proxy server is logging their activity, they may be less inclined to circumvent the policy.

Of course, it's possible that a user has a legitimate need to access unauthorized websites. Organizations typically include a process allowing users to ask for an exception to the policy.

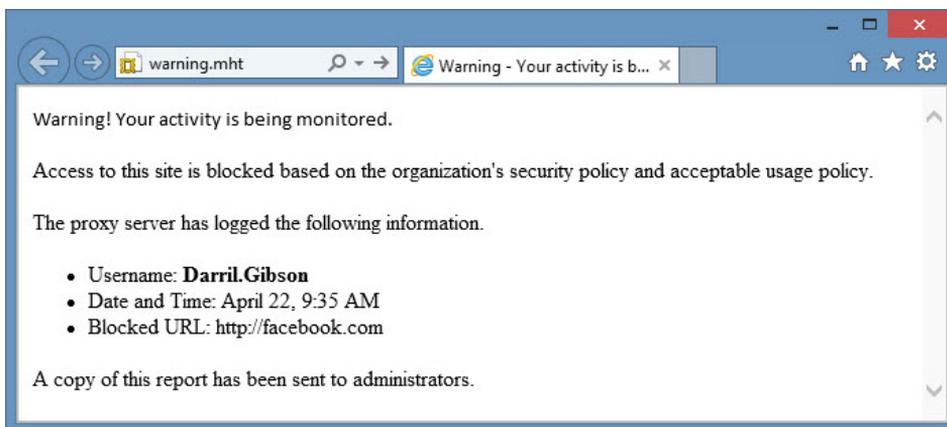


Figure 10-3 Proxy server warning



EXAM TIP Some logs such as proxy server logs can serve as a deterrent to unacceptable or undesirable behavior. Although this isn't a good deterrent to stop criminals, it can help to deter many users from engaging in unsafe computing practices.

Firewall Logs

Most firewalls include the ability to log all activity sent in and out of a firewall. This includes traffic that the firewall blocks, traffic that it passes, or both. These logs can be very useful when trying to reconstruct the events from an attack. For example, after discovering an attack on an internal system, you can use the logs from firewalls at the edge of your network to determine whether the attack came from an external attacker on the Internet or from an internal source.

Intrusion detection systems (IDSs) commonly use firewall logs to detect attacks such as port scans or TCP flood attacks. The IDS regularly monitors the firewall logs, and when it detects a potential attack, it can log and report the event. Active IDSs and intrusion prevention systems (IPSs) can also take action to block an attack. For example, an active IDS can modify the configuration of the firewall to block an attack in progress.

Firewall logs can also provide details on botnet activity. If you suspect that one or more of your internal systems has become a zombie within a botnet, you can check the host-based firewall logs of the system and the network-based firewall logs for suspicious activity. For example, if 20 of your systems are regularly connecting with a system that has an IP address from another country and you don't recognize that address, it could be an indication of botnet-related activity. The foreign IP address might be the address of a command-and-control server used to provide instructions for zombies within the botnet, or another server that attackers hijacked to provide instructions to the zombies.

Reviewing Logs

One of the most important steps of auditing is reviewing audit logs. Reviewers should look for any types of anomalies in the logs. For example, a long string of failed attempts to log on to the administrator account indicates a potential brute-force attack. If the logs then show a successful logon event for the administrator account, it indicates an attacker probably guessed the password for the account and is now accessing the system with administrative privileges. In this situation, it would be important to examine all the activity by this administrator account.

Admittedly, reviewing logs for anomalies isn't easy to do. If you're reviewing thousands of log entries and the log includes one or two abnormal entries, it's easy to overlook the abnormal entries. However, many automated systems are available that can look for and issue alerts for specific events.

In some organizations, administrators get overwhelmed with day-to-day activities and skip the reviewing step. However, at the very least, administrators should archive the logs so that they can review them later, if necessary. For example, Figure 10-4 shows how you

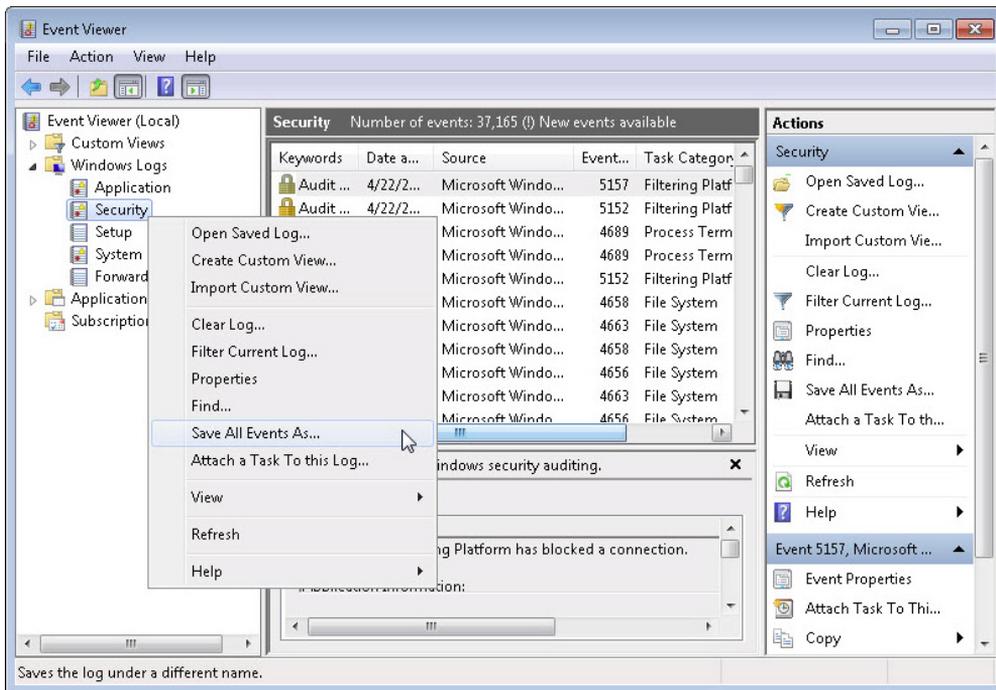


Figure 10-4 Archiving a log

can right-click any log in Microsoft's Event Viewer to save the contents of a log. Administrators and auditors can open saved logs in Event Viewer by selecting Open Saved Log and browsing to the location of the saved file.

When you archive logs like this, you can review them later. In other words, even if you don't notice an incident until much later, you can still re-create the events by reviewing the archived logs.

Managing Audit Logs

In addition to storing logs on remote systems to maintain their integrity, you can take several other steps to improve the security related to logs. A primary step is to protect the logs with permissions.

Access to logs should be restricted to personnel (such as administrators) who have a need to view or maintain the logs. If anyone has the ability to clear the logs, users might erase logs accidentally. Similarly, if anyone can clear the logs, attackers will be able to do so to erase their activity after an attack.



TIP Most systems include the ability to detect when someone has cleared a log. For example, in Windows systems, the first log entry in a new log records details on when someone cleared the previous log. If an organization normally archives and clears logs on Friday, but a log shows that someone cleared it on a different day, it indicates suspicious activity that deserves further investigation.

Many organizations mandate the retention of logs for a specific period such as three years or longer. This allows an investigator to go back in time to reconstruct events even if several months or years have passed. Log retention policies require someone to archive the logs periodically. In smaller organizations, administrators do this manually. Larger organizations use automated processes to archive logs.

Performing Security Audits

A security audit examines an organization's practices and operations to determine whether they conform to the organization's policies or applicable laws. An organization can perform an internal audit to examine its practices, or external auditors can come in to examine the organization's practices. The audit documents the organization's policies, processes, controls, testing, and results. Audits can be performed on a periodic basis, such as once a year, after an event such as a security incident, or when required by laws or guidelines. Audits are detective controls because they detect events after they've occurred.



EXAM TIP Security audits help an organization identify vulnerabilities in its processes and procedures. After a security audit, it's important to implement fixes to mitigate or reduce the discovered vulnerabilities.

For example, an organization can perform monthly vulnerability scans to check for weak passwords. These scans can detect passwords that don't comply with a password policy. Similarly, an organization can perform an annual audit of processes within the organization to determine if these processes comply with the organization's security policy. Organizations also perform these types of audits after an incident to determine if personnel are complying with existing policies. For example, after a successful malware attack, an audit can help determine if the infection was due to personnel not following existing policies.

Organizations that must follow certain laws and guidelines often use compliance audits to ensure that personnel are complying with the relevant laws and guidelines. Internal compliance audits provide proof to outside entities that the organization is taking steps to ensure compliance. For example, if a U.S. organization handles any type of health- or medical-related data, it must comply with the Health Insurance Portability and Accountability Act (HIPAA). If an incident occurs showing a violation of the law, the organization can show documentation from past audits as proof that the organization is taking active steps to be compliant with the law. However, if no past audits exist,

the organization may find itself incurring significantly higher fines and penalties after a security incident.

In other situations, organizations hire external auditors, but typically only when they are required to do so. For example, the Payment Card Industry Data Security Standard (PCI DSS) might require a large organization that handles credit card data to hire an external auditor after a data breach.

Auditing Passwords

A password audit checks to see whether users are following the policies related to passwords. For example, a password policy may dictate that regular users use strong passwords with at least three of the four character types (uppercase, lowercase, numbers, and symbols) and be at least eight characters long. A password audit could use a password-cracking tool to determine whether any passwords fail to meet these requirements. The password audit could also check to see if users are regularly changing their passwords in compliance with a written password policy.

I know from experience that just asking users to use strong passwords and expecting that they will, or just asking them to change their passwords regularly and expecting that they will, works about as well as getting a cat to come to you when you call. In general, you can't expect most users to take secure actions on their own.

Instead, organizations use technical controls to implement written password policies. For example, Microsoft operating systems use a password policy within Group Policy to force users to create strong passwords and change them regularly. Third-party tools are also available that can enforce the use of strong passwords, but they do cost money. However, if a password audit shows that users are using weak passwords and not following the written policy, the cost of a password enforcement tool is justified.

Auditing Security Policies

Many organizations require an annual review of a written security policy to ensure that the policy still meets the needs of the organization. Auditors can examine historical data from the past year to see whether users are following the security policy. They can also examine details of security incidents to help them determine if the policy is actually helping the organization prevent security incidents.

For example, some security policies look great on paper, but employees never hear about them. This often results in users repeating the same mistakes year after year. Social engineers may continue to trick users, or employees might not follow specific procedures simply because they are unaware of the policies. In this case, the audit may discover a lack of awareness and training and recommend specific actions to improve employee awareness.



NOTE Chapter 12 discusses security policies in more depth. A security policy is a written document that provides overall direction for an organization's security.

ISACA

ISACA is an international organization that provides standardization for IT professionals involved in auditing. It includes more than 110,000 members from more than 180 countries working in various IT-related positions, such as IS auditor, internal auditor, consultant, educator, IS security professional, and chief information officer (CIO).

Two of ISACA's top auditing certifications are the Certified Information Systems Auditor (CISA) and the Certified Information Security Manager (CISM) certifications.

- **CISA** The CISA certification identifies individuals who audit, control, monitor, and assess an organization's information technology and business systems.
- **CISM** The CISM certification is management-focused and identifies individuals who design, build, and manage enterprise information security programs.



NOTE ISACA's original name was the Information Systems Audit and Control Association. It dropped the long name in 2008 and now goes by just the acronym ISACA.

ISACA developed the Control Objectives for Information and Related Technology (COBIT) framework. COBIT is a robust framework that helps managers identify processes to manage IT systems and it is one of the popular management frameworks. COBIT 5 is based on five principles and seven enablers, which are generic enough that organizations of any size can apply them.

Exploring PCI DSS Requirements

The Payment Card Industry (PCI) Security Standards Council created and regularly updates the Payment Card Industry Data Security Standard (PCI DSS). The PCI DSS requires that organizations that process credit card payments comply with specific steps to protect customer account data, and it includes a strong focus on protecting credit card data. The goal of PCI DSS is to reduce overall credit card fraud, and toward that goal it specifies 6 control objectives with 12 supporting requirements related to security management, policies, procedures, network architecture, software design, and more.

Table 10-1 provides an overview of the PCI DSS (Version 3.1) objectives along with the specific requirements associated with these objectives.



TIP Although these objectives and requirements only apply to organizations that must comply with PCI DSS, you can slightly modify the steps to apply them to any organization. Replace "cardholder data" with "data" throughout, and this list becomes a generic list for any organization. If you want to read more about the PCI DSS standards, you can access them from here: https://www.pcisecuritystandards.org/security_standards/index.php.

Control Objectives	Requirements
Build and Maintain a Secure Network and Systems	<ol style="list-style-type: none"> 1. Install and maintain a firewall configuration to protect cardholder data. 2. Do not use vendor-supplied defaults for system passwords and other security parameters.
Protect Cardholder Data	<ol style="list-style-type: none"> 1. Protect stored cardholder data. 2. Encrypt transmission of cardholder data across open, public networks.
Maintain a Vulnerability Management Program	<ol style="list-style-type: none"> 1. Protect all systems against malware and regularly update antivirus software or programs. 2. Develop and maintain secure systems and applications.
Implement Strong Access Control Measures	<ol style="list-style-type: none"> 1. Restrict access to cardholder data by business need to know. 2. Identify and authenticate access to system components. 3. Restrict physical access to cardholder data.
Regularly Monitor and Test Networks	<ol style="list-style-type: none"> 1. Track and monitor all access to network resources and cardholder data. 2. Regularly test security systems and processes.
Maintain an Information Security Policy	<ol style="list-style-type: none"> 1. Maintain a policy that addresses information security for all personnel.

Table 10-1 PCI DSS Control Objectives

Compliance Reports

Organizations that fall under the PCI DSS must submit regular compliance reports. These reports identify the business and its use of credit cards. They also document the business's network, including its hardware, software, and connectivity, and how the business complies with each of the 12 requirements. However, organizations do not send the reports to the PCI Security Standards Council. Instead, organizations send the reports to the acquiring bank (the bank handling the credit card transactions) and the global payment brand (such as Visa, MasterCard, and American Express).

Compliance Audits

Businesses are also required to perform periodic audits to verify their compliance with the PCI DSS requirements. Businesses that don't have a high volume of credit card use can perform the compliance audits internally. These organizations can often just answer a Self-Assessment Questionnaire (SAQ) and submit it to the acquiring bank.

However, businesses that have a high volume of credit card use must hire an independent expert to perform an audit. The PCI Security Standards Council certifies Quality Security Assessors (QSAs) and Approved Scanning Vendor (ASV) experts to perform these audits. The PCI Security Standards Council can fine businesses that don't comply with the PCI DSS requirements, or revoke the businesses' authorization to process credit card payments.



NOTE It's worthwhile stressing that the PCI DSS provides a minimum baseline for security. Simply being in compliance does not provide any guarantees against an attacker compromising customer and credit card data. Organizations often implement many other safeguards beyond what is required from the PCI DSS.

Auditing Physical Access Controls

If an organization has physical access controls, it can periodically perform an audit to determine whether these controls are effective. Physical access control audits can also verify if personnel are using the controls or attempting to bypass them.

As an example, many proximity card badge readers have the capability of recording the date and time when an individual enters or exits a secure area. If personnel consistently use their badges to enter and exit the area during the day, the logs would show this. However, if users frequently tailgate each other without using their badges, the logs would show inconsistencies. Each entry showing a user entered the area should have a matching entry for the user when they exit the area.

Chapter 5 covers tailgating in more depth. As a reminder, tailgating occurs when someone passes through a controlled entry point without providing credentials. Instead, the person follows closely behind someone else who has provided credentials. In this scenario, the first person opens the door with a proximity badge and the second person tailgates the first person, bypassing the physical security control.

Physical access logs can also identify whether users are entering secure areas outside of normal work hours. For example, if an employee normally works from 9 to 5 during the week but a log shows the employee entering at 10 P.M. one evening, it might indicate a problem. However, if the logs don't exist, or there isn't a method of reviewing the logs, the organization may not detect this potentially malicious activity.



EXAM TIP A physical access audit can detect whether tailgating is occurring. If employees are able to tailgate each other, a social engineer can use the same technique to get into a secure area.

Understanding Configuration Management

Configuration management is the practice of ensuring that systems are configured with security in mind and that configuration is tightly controlled. Many organizations use configuration management to configure systems with a secure baseline and to ensure the system remains secure with the same configuration.

As an example, administrators ensure that servers start with similar system configurations that make them more secure from the default configuration. Administrators use various hardening techniques such as removing or disabling unnecessary services and protocols, and changing default passwords for all the servers. They then tweak the configuration of different groups of servers to meet the needs for these groups. For example, they configure web servers differently than they configure file servers.

The “Holding Users Accountable with Audit Logs” section earlier in this chapter mentioned FIPS Pub 200, which mandates minimum accounting requirements for IT systems. It also lists the following minimum requirements for configuration control:

- (i) establish and maintain baseline configurations and inventories of organizational information systems (including hardware, software, firmware, and documentation) throughout the respective system development life cycles; and*
- (ii) establish and enforce security configuration settings for information technology products employed in organizational information systems.*



TIP Based on the FIPS Pub 200 requirements to enforce security configuration settings, some federal agencies combine configuration management with change management. However, it’s worth stressing that these are separate processes.

In addition to taking steps to ensure that systems begin with specific configuration settings, an active configuration control plan will mandate that systems be checked periodically. This ensures that the systems maintain the secure settings from the baseline configuration. This process works hand in hand with a strong change management program that attempts to prevent unauthorized changes and the unintended outages caused by unauthorized changes.

Using Imaging for Configuration Management

A common configuration management practice is to use imaging technologies. For example, an organization can deploy each desktop system with a common image that includes the operating system, basic applications, and specific settings needed for usability and security. Images help ensure that every deployed system starts with a secure baseline.

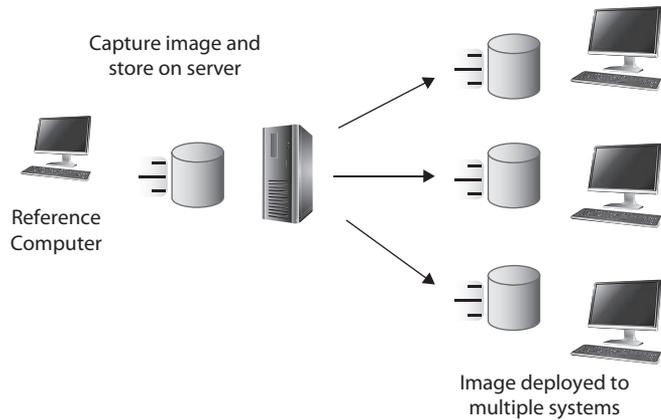


EXAM TIP System baselines are an important element of configuration control and are often implemented with images. When administrators know the starting point (the baseline configuration), they can easily check systems to determine if the configuration is different from the baseline.

Figure 10-5 shows the overall process of imaging. You start with a clean computer called the reference computer and install an operating system on it. You can then install applications and configure the operating system based on the needs of the organization. For example, you can configure security settings such as auditing, a password policy, an account lockout policy, a password-protected screen saver, and much more. After testing the image, administrators capture it using imaging software. They then deploy the image to other systems, and these imaged systems start with the same configuration as the reference computer.

As an example, the U.S. government uses the United States Government Configuration Baseline (USGCB) guidelines to create images for Windows 7 and Red Hat Enterprise

Figure 10-5
Using imaging
for configuration
control



Linux 5 systems. The U.S. Office of Management and Budget (OMB) mandates the use of similar images, and government IT personnel have steadily improved them over the years. These images include over 300 settings that configure and secure systems.

Experts derived the USGCB image settings from the Federal Desktop Core Configuration (FDCC) image. Before the FDCC image was available, individual agencies were responsible for deploying their computers in a secure manner, but many smaller agencies didn't have the knowledge and expertise to secure these computers. They would often repeat the same basic mistakes that left many systems vulnerable in the past. A group of IT security professionals identified the secure settings for the FDCC and created the image, and now agencies use similar images to ensure that computers start with a secure baseline, no matter what level of IT security expertise the agency has within it.

Images also make systems easier to maintain. Help desk professionals only need to learn the configuration of one system. No matter what user they are helping, the configuration is the same, enabling help desk professionals to troubleshoot and restore the system quicker.

In contrast, when an organization has 100 computers and each of these computers has a different configuration, support personnel must first figure out how the system should work normally and then take steps to restore it to its original configuration. Effective configuration management programs reduce the overall total cost of ownership of systems by reducing the repair times.

Using Group Policy for Configuration Management

In Microsoft networks, administrators use Group Policy to configure systems. Group Policy allows centralized administration of systems. An administrator can configure one setting with a Group Policy object (such as opening a port on a firewall), and Group Policy applies the change to all the computers within the scope of the Group Policy object. Whether the organization has 50 computers or 50,000 computers, Group Policy can apply the change consistently to all the computers in the organization.

Group Policy also supports targeting specific systems, such as all the computers in a specific department or only computers running Windows 7. Administrators with a little knowledge of Group Policy can easily implement configuration control to all their Microsoft systems with very little administrative effort.

Understanding Change Management

Many organizations have a change control process in place. The primary purpose is to ensure that changes don't result in unintended outages. Networks are complex, and many of the systems within a network interact with each other. Making a change in one system can easily affect other systems and result in an unintended outage. However, if experts examine these changes before technicians or administrators implement them, the organization can prevent many of the unintended outages from changes.

For example, imagine that a server hosts an application that requires opening TCP port 4567 on a server. Administrators open the port when they configure the server application, and the application runs as expected. Later, a security professional performs a vulnerability assessment and notices TCP port 4567 is open. Thinking this port is a vulnerability, the security professional decides to close the port and reduce the server's attack surface. Of course, this action breaks the server application. Users begin to complain and the organization calls in the application developers to troubleshoot the problem. However, it might take them hours, or even days, before they discover the closed port.

Even though the security professional's intentions are good, the action still caused an outage. You could fill volumes with similar examples of well-intentioned IT personnel making changes, often to correct a minor problem, and in the process creating other problems that cause outages.

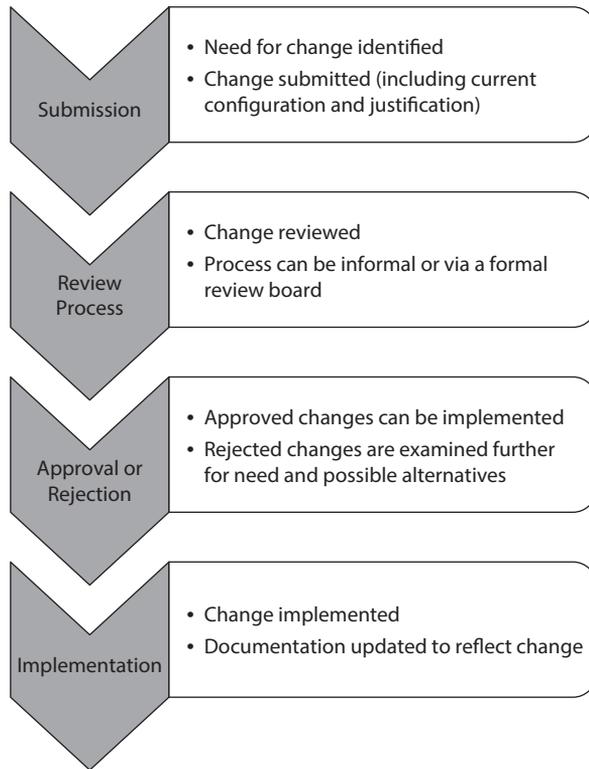
In contrast, a change control process ensures that appropriate personnel examine all changes before approving them. It also provides a documentation trail. If the server ever suffers a catastrophic failure, administrators can restore it to its last state, including all the changes applied to the server.

Although many organizations have adopted some level of change management, there is a wide diversity of maturity among organizations. For example, some organizations have a policy for change management, but follow it only for major changes and still suffer outages from many minor changes.

A change control process gives stakeholders a formal mechanism to propose changes to a system. Multiple entities then examine the change request and either approve or reject it. The following list describes the overall steps in a change control process, as shown in Figure 10-6:

- 1. Change submitted** Someone identifies the need for the change and submits a change request. The change request identifies the current configuration and the desired change, along with justification for the change.
- 2. Change reviewed** The review process allows experts within the organization an opportunity to review the proposed change. For simple changes, this can be a quick, informal process, but for major or potential problematic changes, this

Figure 10-6
Change control
process steps



can require a thorough review by a formal change review board (CRB) or change authorization board (CAB).

- 3. Changed approved or rejected** The CRB or CAB either approves or rejects the change request. They also analyze rejected changes to see whether a change is required, and if so, they look for suitable alternatives.
- 4. Change implemented** Personnel implement and document the change. It is also important to ensure that appropriate personnel know about the approved change. If administrators don't know about changes, they might notice them and then take steps to undo them.

As an example of a change management process, one organization where I worked used an intranet website to submit, review, and approve changes. Experts in different departments reviewed these changes and either approved or suspended them. If each expert approved the request, the CRB approved the request. However, if any expert suspended the request, the CRB reviewed the request at the next meeting.



EXAM TIP A change control process allows stakeholders to propose changes and gives the organization an opportunity to review the changes before implementation. The goal is to reduce unintended outages from unauthorized changes.

Chapter Review

Auditing is an important component of enforcing accountability. When an organization implements strong authentication and authorization methods, administrators and security personnel can use auditing methods to track user activity and hold individuals accountable for their activity. Most auditing methods use clipping levels to generate alerts only when the system detects the number of events exceeds a predetermined threshold. The auditing system ignores events that don't meet the threshold. Together, logs create an audit trail that administrators and security professionals can use to reconstruct the events that occurred leading up to and during an incident.

Many different types of audit logs exist. Operating system logs record events on individual systems. These include security logs, system logs, and application logs. Security logs record security events, such as when a user accesses a file. System logs record system events, such as when a service stops or starts. Application logs record application events based on the needs of the application. Proxy logs and firewall logs record activity through proxy servers and firewalls, respectively. An important step related to logs is to review them on a regular basis.

Security audits examine an organization's practices and operations to determine whether they conform to existing policies or applicable laws. Audits are often required to ensure that an organization is complying with specific external requirements (such as PCI DSS) or laws (such as HIPAA), and internal audits show an attempt to comply with these external requirements. In some situations, organizations are required to hire external auditors to perform audits.

Configuration control ensures that systems are configured using a baseline. This ensures that systems are configured in a secure manner and that configuration is similar (if not identical) for similar systems. Imaging is a common method used to deploy systems with similar needs and configuration requirements. Change management helps prevent unauthorized changes resulting in unintended outages.

Questions

1. Your organization uses strong authentication and authorization mechanisms and has robust logging capabilities. Combined, what do these three elements provide?
 - A. Guaranteed security
 - B. Prevention of unintended outages from unauthorized changes
 - C. Accountability
 - D. Configuration control

2. An accounting system ignores logon failures until an account has three logon failures within a 30-minute period. It then generates an alert. What is the accounting system using?
 - A. Account lockout
 - B. Password policy
 - C. Snipping level
 - D. Clipping level
3. A system ignores potential security violations until it detects a specific number of events. It then raises an alert. What does this describe?
 - A. Clipping level
 - B. Acceptance level
 - C. Audit level
 - D. Baseline level
4. Which of the following statements best describes a benefit of using clipping levels?
 - A. Clipping levels ignore baselines and generate alerts when they detect security violations.
 - B. Clipping levels ignore normal user errors, but generate alerts when these errors exceed a predetermined threshold.
 - C. Audit trails use clipping levels to record all potential alerts for accountability.
 - D. Clipping levels ensure systems generate alerts when they detect any potential security violations.
5. A user entered an incorrect password three times. Now, the user is no longer able to log on. What caused this to occur?
 - A. Password policy
 - B. Account lockout policy
 - C. Clipping level
 - D. Audit trail
6. What do you call a group of one or more logs used to re-create events leading up to and occurring during an incident?
 - A. A configuration control program
 - B. A change management program
 - C. A security audit
 - D. An audit trail

7. What type of control is an audit trail?
 - A. Preventive control
 - B. Corrective control
 - C. Detective control
 - D. Physical access control
8. What type of log on a Microsoft system records auditable events, such as when a user deletes a file?
 - A. System
 - B. Security
 - C. Application
 - D. Forwarded Events
9. Of the following choices, what is an example of an auditable event logged in an operating system's security log?
 - A. Access through a firewall
 - B. Accessing a website through a proxy server
 - C. Reading a file
 - D. The date and time when a service starts
10. Of the following choices, what is the best example of a log used as a deterrent for internal employees?
 - A. Proxy server log
 - B. Network firewall log
 - C. Security audit
 - D. Change management log
11. You suspect that many internal systems may be part of a botnet. What log would you review to verify your suspicions?
 - A. Network-based firewall logs
 - B. Host-based firewall logs
 - C. Operating system logs
 - D. System security logs
12. Of the following choices, which one is NOT a recommended strategy for audit logs?
 - A. Review the logs regularly
 - B. Archive logs for later review
 - C. Periodically overwrite logs
 - D. Store logs on remote servers

13. What is the purpose of reviewing logs?
 - A. Detecting potential security events
 - B. Preventing potential security events
 - C. Correcting potential security events
 - D. Deterring potential security events
14. Your organization has recently completed a security audit. Which of the following is NOT a valid step to take after completing the audit?
 - A. Approve changes
 - B. Evaluate controls
 - C. Implement fixes
 - D. Delete the security audit
15. An organization handles credit card data from customers on a regular basis. What provides the security objectives and requirements that the organization must follow?
 - A. PCI DSS
 - B. HIPAA
 - C. FIPS Pub 200
 - D. NIST SP 800-53
16. A badge reader records employee names, dates, and times when employees enter and exit a secure server room. An auditor reviewed the logs and noticed that they showed that many employees entered the room, but the logs do not show when all of the employees exited the room. What does this indicate?
 - A. The badge reader is operational
 - B. Tailgating
 - C. The mantrap is not being used
 - D. Unauthorized entry
17. Who would measure the effectiveness of an organization's security controls?
 - A. An administrator
 - B. A manager
 - C. An auditor
 - D. A data owner
18. Of the following choices, what is a primary method used for configuration control?
 - A. Baseline
 - B. Change management requests
 - C. Security logs
 - D. Password audits

19. Which of the following is NOT a valid method used for configuration control?
- A. Imaging
 - B. Microsoft's Group Policy
 - C. Change management
 - D. Proxy server logs
20. Of the following choices, what can help ensure that system modifications do NOT cause unintended outages?
- A. Security audit
 - B. Change management
 - C. Configuration control
 - D. Audit trail

Answers

1. **C.** Authentication, authorization, and accounting (AAA) provide accountability, and logging provides the accounting element. Although AAA increases security, it does not guarantee security. Change management prevents unintended outages from unauthorized changes. Configuration control ensures systems are deployed with a secure baseline and maintain approved configuration settings for system stability.
2. **D.** A clipping level uses a predetermined level as a threshold. A classic example is three or five logon failures in a short period, such as within 30 minutes. Although many operating systems use account lockout policies to actually lock the account after a predetermined level, the question doesn't ask what happens to the account, but instead asks what the accounting system is using to ignore the first two logon failures and only generate the alert after three logon failures. A password policy ensures that users have strong passwords and change them regularly. Snipping level isn't a valid term associated with accounting systems.
3. **A.** A clipping level uses a predetermined number of events as a threshold. An auditing system ignores the events until it detects the number of events has exceeded the threshold level. Acceptance level and audit level are not valid terms. The question doesn't describe a baseline level.
4. **B.** Clipping levels ignore normal user errors, but generate alerts when these errors exceed a predetermined threshold. Clipping levels are not associated with baselines. It is possible to configure an audit trail without clipping levels, but if clipping levels exist, the audit trail does not ignore them. Clipping levels do not generate alerts when they detect any potential errors or security violations, but instead only generate alerts when they detect the number of events has exceeded a predetermined threshold.

5. **B.** An account lockout policy locks out an account after a predetermined number of failed logins. The password policy ensures that users create strong passwords and change them often. The account lockout policy is using a clipping level by ignoring failed login attempts until it detects a preset threshold, but the clipping level doesn't lock the account. An audit trail is one or more logs used to reconstruct events leading up to and occurring during an incident.
6. **D.** An audit trail is one or more logs that can re-create events leading up to and occurring during an incident. Configuration control helps ensure that systems are configured in a secure manner. A change management program allows stakeholders to request changes and helps reduce unintended outages from unauthorized changes. A security audit examines an organization's policies and procedures to determine whether the organization follows these policies and procedures.
7. **C.** An audit trail is a technical detective control, because it uses technology and can detect incidents after they occur. A preventive control attempts to prevent incidents. A corrective control attempts to reverse the impact of an incident after it has occurred. A physical access control is an item that you can physically touch.
8. **B.** The Security log records auditable events, such as when a user accesses or deletes a file (as long as resource auditing is enabled). The System log records system events such as when a service stops or starts. The Application log records application events. The Forwarded Events log shows events forwarded from other systems as part of an event subscription.
9. **C.** A security log records auditable events related to resources, such as when a user reads, modifies, or deletes a file. Firewall and proxy server logs are not operating system logs. A system log would record events such as when a service stops or starts, but not security events.
10. **A.** A proxy server log can serve as a deterrent for internal employees. If employees know that the server is monitoring and logging their activity, they may be less likely to engage in activity that violates the security policy. A network firewall log can capture activity for traffic to and from the Internet, but does not provide much of a deterrent for internal employees. A security audit is not a log. A change management log documents changes for a system.
11. **A.** Network-based firewall logs record traffic on the network, and because many systems are involved, network-based firewalls is the best choice. Each of the other logs are local logs on individual systems. This would require checking logs on multiple systems, rather than checking logs on a single network-based firewall.
12. **C.** If you periodically overwrite logs, it is no longer possible to review the logs. However, all of the other choices (review the logs, archive the logs, and store logs on remote servers) are recommended strategies to retain the integrity of audit logs.
13. **A.** Security professionals and auditors can detect potential security events by reviewing logs after the event has occurred. Reviewing the logs doesn't prevent an incident that has already occurred, and reviewing the logs does not enable security professionals and auditors to correct the effects of an incident. While

logging some activity, such as proxy servers, can deter events, reviewing the logs doesn't deter the activity.

14. **D.** Organizations should keep security audits instead of deleting them. This allows personnel to use them as reference points in future audits. Audits typically evaluate existing controls and recommend changes. Management approves the changes and directs personnel within the organization to implement the fixes.
15. **A.** The Payment Card Industry Data Security Standard (PCI DSS) provides 6 control objectives and 12 supporting requirements that organizations must follow if they process credit card payments from customers. The Health Insurance Portability and Accountability Act (HIPAA) covers organizations handling health- and medical-related data. Federal Information Processing Standard Publication 200 (FIPS Pub 200) identifies standards required by federal agencies. NIST SP 800-53 provides information on recommended security controls.
16. **B.** Logs that include entries showing employees entered a secure area but do not include entries showing they exited indicate tailgating is occurring. Some employees are using their credentials to exit (and the logs show them exiting), but other employees are following closely behind these employees without showing their credentials (and the log doesn't include entries for these employees). While it is possible the badge reader has a problem, it is recording some employees exiting, so this isn't the most likely cause. Mantraps prevent tailgating, and if a mantrap is in use, employees would be forced to use it. There isn't any indication of unauthorized entry.
17. **C.** An auditor would measure the effectiveness of a security control. An internal auditor might have other roles, such as an administrator, a manager, or a data owner. However, when measuring the effectiveness of security controls, they are acting as an auditor.
18. **A.** A baseline is a primary method used for configuration control and it ensures that systems start in a known state. Automated or manual processes periodically examine the systems to verify the system still has the same configuration settings from the baseline. An organization doesn't approve and implement all change management requests, so examining the requests does not give an accurate representation of the server configuration. Security logs and password audits aren't typically used for configuration control.
19. **D.** Proxy server logs record what websites users visit and are not used for configuration control. Each of the other choices can be used for configuration control.
20. **B.** A change management program allows stakeholders to request changes and helps reduce unintended outages from unauthorized changes. A security audit examines an organization's policies and procedures to determine whether those who work in the organization follow these policies and procedures. Configuration control helps ensure that systems are configured in a secure manner and similarly to each other. An audit trail is one or more logs that can re-create events leading up to and occurring during an incident.

This page intentionally left blank

Security Operations

In this chapter, you will learn about

- Data classifications such as Confidential, Sensitive, Private, and Public
- The difference between data at rest and data in motion
- Different elements of data management policies
- Components of a database such as tuples, rows, primary keys, and foreign keys
- Risks related to data inference, data diddling, and Big Data
- Regulatory requirements related to PII and PHI
- Asset management components
- The differences between certification and accreditation
- The Common Criteria and its evaluation levels
- Using a risk management framework with a certification and accreditation process
- Different phases of a system development life cycle

Handling Data

Data is one of the most valuable resources an organization has, so it's important to understand the value of data and some of the risks related to data. Different types of data have different value, so an organization protects data based on its value. Highly classified or sensitive data requires a higher level of protection. Similarly, when an organization no longer needs particular data, it uses one of several different methods to destroy the data depending on its value.

Classifying Data

One of the first steps in protecting data is identifying what data needs protection by classifying it. For example, an organization will often take additional steps to encrypt and protect highly classified data, such as research and development data that took years to develop. However, data published in company brochures or company websites does not deserve the same level of protection. By classifying data, it becomes easier for anyone within the organization to identify its worth to the organization and the level of protection it deserves.

For example, the U.S. government classifies data as Top Secret, Secret, Confidential, and Unclassified. The higher the classification, the more protection it requires. The following list shows the definitions assigned to each of these classifications:

- **Top Secret** The public disclosure of Top Secret material could cause exceptionally grave damage to U.S. national security.
- **Secret** The public disclosure of Secret material could cause serious damage to U.S. national security.
- **Confidential** The public disclosure of Confidential material could cause damage to U.S. national security.
- **Unclassified** The public disclosure of Unclassified material will not cause damage to U.S. national security. People can obtain information in this category through the Freedom of Information Act.

To have access to any of this classified material, an individual must have the associated clearance. For example, someone with a Secret clearance may be able to view Secret material, but not Top Secret material.



TIP The principle of need-to-know is widely used to protect data. For example, just because individuals have a Secret clearance doesn't mean they can access all Secret material. The clearance level only indicates they are cleared to view the material if their job requires it, but does not grant access. If their job does require access to the material, an administrator grants them appropriate access.

Many private organizations also classify data, although they rarely use the same Top Secret, Secret, Confidential, and Unclassified labels used by the U.S. government. An important point to remember is that organizations can use any labels they desire. For example, an organization can use Class 0, Class 1, Class 2, and Class 3. However, end users might find them to be a little mysterious. What's the highest classification, Class 0 or Class 3? One organization might say Class 0 is highest, and another organization might say Class 3 is highest. Of course, both organizations are correct, because the organization defines the labels in a security policy or data policy.

Data owners are responsible for identifying data classifications. In civilian organizations, data owners are typically the chief executive officer (CEO), chief information officer (CIO), or department heads. A data owner ensures that a security policy or data policy includes the labels and accompanying descriptions.

Here are some examples of data classifications from highest classifications to lowest:

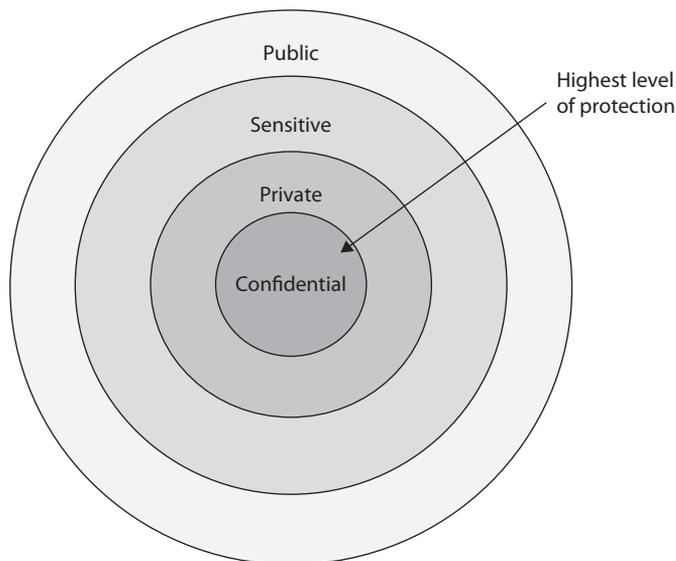
- Confidential, Private, Sensitive, Public
- Proprietary, Internal Use Only, Sensitive, Public
- Private Highly Restricted, Private Restricted, Private, Unrestricted
- Highly Confidential, Restricted, Private, Public

- Regulatory, Confidential, Internal Use Only, Public
- Confidential, Restricted, Internal Use Only, Public
- Highly Confidential, Restricted, Private, Unrestricted

Figure 11-1 shows the Confidential, Private, Sensitive, and Public labels that any organization can use. It also includes possible definitions that an organization can use within its security policy or data policy. Notice that Confidential data is at the core and deserves the highest level of protection.

- **Confidential** This includes proprietary information central to the operation of a company. It could include research and development information for upcoming projects, trade secrets for products sold by the company, and so forth. If someone discloses an organization's confidential data, it could result in an exceptionally grave negative impact on the mission of the organization.
- **Private** This includes information used privately within the company, such as employee records and customer data. Private data is intended for internal use only, and if it is disclosed, it could result in a serious negative impact on the mission of the organization.
- **Sensitive** This includes information that requires special precautions for its protection, such as financial information about a company. If someone discloses sensitive data, it could result in a negative impact for the organization.
- **Public** This is data that is either publicly available or would not cause any harm if it was publicly available. If disclosed, this data would not have any negative impact on the organization.

Figure 11-1
Data classifica-
tions used
by some
organizations





TIP Organizations aren't required to use any specific data labels, so it's entirely possible an organization could have different labels or define these a little differently. The important thing to grasp is that organizations do classify material and provide higher levels of protection for data that requires it.

Marking and Labeling Data

After identifying data classifications, personnel mark or label the data. In general, *marking* refers to using digital methods to identify the data classification. *Labeling* refers to using physical labels to identify the data classification. This ensures that users can immediately identify the classification files and media holding the files. For example, after printing classified data, users may be required to place it within a folder with a cover sheet indicating the classification of the data. Digital files can include watermarks, headers, or footers that show the classification so that it is easily identifiable after a user prints it.

It's also important to label hardware that holds or processes classified information. This includes portable media such as USB drives and tapes, cabinets where portable media is stored, and computer systems that process classified data. Organizations commonly use color-coded stickers with the classification name. As an example, they could use red stickers with the label Proprietary. This allows employees to identify the classification based on the name and the color of the label.

An alternative is to purchase a large quantity of drives (such as red USB drives) and train employees to use only these drives for classified data. It's also possible for administrators to configure technical policies to encrypt data on these drives automatically. These technical policies identify the drives using a universally unique identifier (UUID).

Roles and Responsibilities

Multiple people within an organization have different responsibilities for how data is handled, and those responsibilities can be classified by their role. Some of the common roles and responsibilities are

- **Management** Personnel in management and leadership positions define the data classifications and specify the requirements to protect data with different classifications. A data policy or security policy includes these definitions, and management personnel ensure the policy is available to everyone within the organization.
- **Data owner** Data owners have primary responsibility for protecting the data based on its classification and requirements stated within a security policy. When data owners create data, they identify the appropriate classification and mark or label the data according to requirements dictated in the data policy. They periodically review the data and modify the classification if necessary based on changes within the organization.

- **Custodian** In some cases, the data owner delegates various data-related tasks to a data custodian. For example, a data custodian could perform regular backups of the data based on requirements stated within a backup policy.
- **Administrator** Administrators use available access control mechanisms to provide access to the data to personnel that need access. For example, administrators assign permissions to personnel to grant access, typically by direction from the data owner.
- **User** Users access the data. Users are also responsible for protecting the data they handle.



TIP Data owners are ultimately responsible for protecting the data. This includes ensuring that it is marked or labeled correctly and that security controls match the requirements of the data classification.

Protecting Data from Cradle to Grave

Personnel protect classified data from the moment someone creates it until someone destroys it. The level of protection is based on the classification. If classified data is stored or processed on a computer system, then the system must be protected at least at the same level as the classified data. Similarly, if classified data is transmitted over a network, all access points on the network must also be protected to at least the same level.

If a system or hard drive fails, it's important to remove all remnants of data before disposing of it. For example, just because a hard drive fails during normal operation, that doesn't mean that all of the data on the drive is inaccessible. Experts working in a clean room can open failed disks and often retrieve data remnants. Although this process is expensive, it can be worthwhile if it prevents an attacker from retrieving sensitive data from an organization.



TIP The "Removing Data Remnants" section later in this chapter covers various methods of removing residual data.

Data at Rest and Data in Motion

Data at rest is any data that is in computer storage, such as on system hard drives, portable USB drives, flash drives, storage area networks, backup tapes, and so on. Data in motion (sometimes called data in transit) is any data transmitted over a network. This includes data transmitted over an internal network using wired or wireless methods and data transmitted over public networks such as the Internet.

Encryption is often used to protect data at rest and data in motion. Encryption converts plain-text data into cipher text, making it unreadable until it is decrypted. Strong encryption algorithms ensure that unauthorized entities are unable to read encrypted data.



EXAM TIP Encryption protocols such as Advanced Encryption Standard (AES) encrypt data at rest. Protocols such as Secure Sockets Layer (SSL), Transport Layer Security (TLS), and Internet Protocol security (IPsec) protect data in motion. Chapter 14 covers these and other encryption protocols.

Additionally, access control methods such as strong authentication and authorization practices help to protect data. Access controls help ensure that only authorized people have access to a network where data is stored and transmitted. However, access controls aren't useful when transmitting data over a public network.

It's also important to protect data during active processing (data in use), and this is a function of the application processing the data. For example, imagine a web application retrieves customer data from a database, including credit card data, to process a purchase. Application developers often use specific buffers for sensitive data and purge these buffers when the data is no longer needed. In this example, the data goes through several stages:

- **Data at rest** Sensitive data is stored in an encrypted format within the database. However, the database decrypts it before sending it to the web application.
- **Data in motion** The web server and database server create a secure channel to transfer the data. For example, they can use TLS to encrypt the sensitive data while transmitting it over the network.
- **Data in use** The web server decrypts the sensitive data and the web application processes it in an unencrypted format. When the web application is done with the data, it clears the data buffers holding this data.

Data Management Policies

Many organizations implement data management policies to help employees understand the value of different types of data. These policies also identify methods of protecting sensitive data. These policies help an organization protect data from unauthorized disclosure and misuse throughout the lifetime of the data. The following subsections identify some of the common elements of data management policies.

Storage Media

Storage media refers to how data at rest is stored. It can refer to data on hard drives, solid state drives (SSDs), portable USB drives, USB flash drives, optical disks, backup tapes, or any other media storing data. Data management policies may restrict the use of some media and require special handling of any media that holds sensitive data.

One of the biggest risks of most storage media is that it's highly portable. It's very easy for someone to copy mass amounts of data onto the media, and the organization can lose control of it. For example, a malicious insider can steal valuable data by copying it onto a portable USB drive. It's also possible for an honest employee to lose a USB flash drive after copying data onto it.

A common way to protect the confidentiality of sensitive data on storage media is with encryption. For example, many USB drives are available that automatically encrypt

data when users copy data to them. Several technical security controls allow administrators to configure whole disk encryption for portable devices, and hard drives within computer systems.

In addition to encryption, the media should be labeled to indicate the highest classification of data stored on the media. When the media is not in use, employees should secure it in a storage cabinet or safe.

Transmission

Data transmitted over a network is data in motion. Attackers can use sniffers or protocol analyzers (covered in Chapter 5) to capture and read the data transmitted in cleartext. However, encrypting sensitive data before transmitting it thwarts a sniffing attack. Data management policies dictate the minimum encryption methods of specific types of data.

Chapter 3 introduced several encryption protocols that encrypt data before transmission. Secure Shell (SSH) is often used to encrypt traffic such as Telnet and File Transport Protocol (FTP) on internal networks. IPsec can be used to encrypt traffic in internal networks and traffic sent over the Internet, such as with virtual private networks (VPNs). SSL and TLS are two protocols commonly used to encrypt data transmitted over the Internet, such as with HyperText Transport Protocol Secure (HTTPS).

Archiving and Retention Requirements

Archiving refers to making a backup of data and keeping it for long-term storage. Chapter 9 covers backup strategies in more depth. *Retention* refers to how long archives and other data are kept. Data policies often identify both archiving and retention requirements.

Laws or regulations may require an organization to keep data for a minimum time-frame, and sometimes indefinitely. For example, laws in the United States require that e-mails sent to and from officials in several U.S. public offices be kept forever. These archives are kept for historical purposes. Similarly, many banking regulations require institutions to keep certain data for three, five, or seven years.

Retention requirements also identify how long to retain backups. For example, an organization can decide to keep copies of backups for the last three years or the last 30 days. Many factors go into these decisions, and an important one is cost. The number of backup tapes required to keep archive copies of backups for the last three years can be quite high.

Additionally, many companies consider legal consequences of long-term storage of data. For example, a court order might direct a company to turn over all e-mail sent or received by a certain individual over the last five years. However, if the company retention policy states that e-mail is only kept for a year, the company only needs to retrieve e-mails from the last year. This assumes that personnel are following the policy and the company only has a year's worth of e-mails. In contrast, if the organization retains e-mail for five years, IT personnel will have to sift through five years of e-mails to comply with the court order. This can be time consuming and labor intensive.

Removing Data Remnants

Data remnants refers to the residual data left on media. Data management policies usually include information on how to eliminate data remnants. This allows personnel to

reuse the media or safely dispose of the media without losing confidentiality of the data. Methods of eliminating data remnants include sanitizing (or purging) the media and destroying it.



TIP Residual data is sometimes referred to as data remanence. However, remanence technically refers to residual magnetic flux on traditional hard drives. It does not apply to other media such as CDs, DVDs, and solid state drives.

- **Sanitizing (or purging)** Sanitization methods remove all residual elements of data, preventing someone from accessing it. After sanitizing media, personnel can reuse it or dispose of it.
- **Destruction** In some cases, sanitizing methods are not reliable enough, so personnel destroy the media instead. For example, if media holds highly classified data, an organization's data policy may require personnel to destroy the media instead of using methods to purge it of all data remnants.

If employees throw paper documents into the trash, dumpster divers can retrieve them. Instead, employees should destroy paper documents either by passing them through a crosscut shredder or by burning them in an incinerator. It's important to use a crosscut shredder because it cuts the paper into smaller pieces that resemble confetti. In contrast, strip-cut shredders cut the paper into long strips, which a patient attacker can reconstruct more easily.

Most administrators understand that simply deleting a file or formatting a drive doesn't actually remove the data. Someone with the right software and a little bit of expertise can often retrieve the data. To combat this, organizations implement different procedures to remove data. In some cases, an organization approves a specific application to write a pattern of ones and zeros to remove data. These programs can "shred" a single file by overwriting it, or completely erase the entire disk by overwriting it.

Magnetic-based media such as traditional hard drives and backup tapes store data by magnetizing elements of the media. *Degaussing* uses a powerful magnet to remove data and is an effective method of sanitizing backup tapes. You can also degauss hard drives, and this usually destroys the drives. Optical media such as DVDs and CDs do not use magnetic methods, so degaussing does not remove data from this type of media. Similarly, degaussing won't remove data from an SSD because SSDs use integrated circuitry instead of the magnetic flux used on hard drives. Instead, optical media and SSDs must be physically destroyed.

SSDs present a serious risk to an organization when it is sanitizing systems. If you put a traditional hard drive next to an SSD, they look the same. Personnel might think that they can degauss both and then throw them away. However, the SSD will still hold all the data.

When destroying media, the goal is to ensure it is unusable, which is the most secure method of protecting confidentiality of data. Destruction methods include shredding it (for media such as optical media) and burning or incinerating it (for backup tapes). Policies

often require personnel to open hard drives and destroy the platters. For example, personnel can remove the platters and send them through large shredders or incinerate them.



TIP Data policies identify how to destroy data and eliminate data remnants. It's common to degauss backup tapes to remove data. Programs that repeatedly write a series of ones and zeros onto a drive overwrite or shred files.

Cloud Computing

As an example of a data management policy, the U.S. Department of Defense (DoD) released its *Cloud Computing Security Requirements Guide* in January 2015. It identifies specific requirements for storing data in the cloud.

The guide includes six information impact levels. It defines data within each level and defines the security controls required when storing the data in the cloud. For example, Level 4 is Controlled Unclassified Information such as personally identifiable information (PII) or protected health information (PHI), and Level 6 includes any classified information up to Secret.

Requirements to store the data increase for higher levels. For example, the list of requirements for storing Secret Level 6 data in the cloud is quite extensive and includes several elements such as legal jurisdiction and location requirements. Secret Level 6 data must be located in a facility under the legal jurisdiction of the United States.

The guide also identifies specific requirements for encryption of data, destruction of data, and disposal of storage media. For example, personnel can reuse media holding Level 4 data after purging the media using an acceptable method. In contrast, media holding Level 6 data cannot be reused to hold anything other than Level 6 data.

Deduplication

It's common for multiple individuals to have a copy of the same file. If this file is stored in multiple locations, it can take up a significant amount of wasted storage space. *Deduplication* refers to ensuring that a file is stored only once on a system, even if multiple users have access to the same file.

For example, employees commonly have private shares on servers where they can store their data. Imagine if 100 employees have shares on the same server, and each employee has a copy of a 5MB company policy document. This single file takes up 500MB of space if there are 100 separate copies of it. In contrast, with deduplication the file will be stored only once and take up only 5MB of storage. Each employee will appear to have a copy of the document in their private share. However, the file that the employees see acts as a pointer to the single document.

Additionally, many deduplication systems support multiple versions of the same file without keeping multiple copies of the file. Consider the same 5MB file. If five users make slight modifications to their copy, a deduplication system can record their changes in separate files. The system keeps the original 5MB file intact. If any of these five users

opens the file, the system applies the changes that the specific user has made to the file. This is all automatic and often done without the user's knowledge.

There is a trade-off between deduplication and encryption. If the files are encrypted (and the deduplication software cannot decrypt them), then the deduplication software cannot identify any two files that are the same. While there are methods of weakening the encryption so that the deduplication software can identify identical files, this sacrifices security. Any method of weakening the encryption increases the possibility that an attacker can decrypt the data. With this in mind, unless you weaken security (certainly not recommended), you'll rarely see gains from deduplication software.

Data Loss Prevention

Data loss prevention (DLP) systems attempt to monitor data usage and prevent the unauthorized use or transmission of sensitive data. Different types of DLP systems can monitor data in motion and data at rest.

Network-based DLP devices have the ability to scan all traffic that exits a network. This includes e-mail and attachments. These devices can look for specific data and block it before it leaves the network.

For example, a company may want to ensure that e-mail or attachments never include Social Security numbers. A Social Security number is in the format of *xxx-xx-xxxx*. The device can scan all outgoing traffic looking for this match. When it detects a match, it blocks the traffic from leaving the network and sends an alert to an administrator. Similarly, if a company uses classification labels such as Confidential, Private, and Sensitive, it can configure the device to look for these words along with any other words or patterns of interest.

Some DLP solutions can limit the use of hardware. For example, an organization may want to prevent users from copying certain data to external USB drives or printing this data to a printer. The DLP software scans the data before copying it or sending it to the printer. If it includes restricted data, such as Social Security numbers or specific classification labels, the DLP software blocks it.

Software DLP methods have the ability to scan data stored on a system to detect inappropriate data. As a simple example, an application can scan data stored on a shared server, looking for MP3 files or other data types that the company has stated should not be stored there. These applications can monitor all data stored on a system and periodically provide reports to administrators.

DLP systems can usually perform deep inspection of the data. For example, even if several files are embedded in a ZIP file or other archive file type, the DLP system can still extract the data to inspect it.



EXAM TIP Data loss prevention (DLP) systems help prevent data from leaving a network. A network-based DLP can scan all traffic looking for specific data and block transmissions when it detects restricted data transmissions.

Social Network Usage

Social networks also present risks related to the loss of data. These include sites such as Facebook, Twitter, and LinkedIn.

Chapter 5 mentions some of the risks with social networking sites in the context of social engineering and phishing. Attackers frequently send phishing messages to users of social networking sites, attempting to trick them into following a link to a malicious website or downloading malicious software. However, there are additional risks associated with data from social networking sites.

For example, a chain letter was floating around asking people on Facebook to list ten things about themselves that others don't know. It included questions about the person's favorite things (like favorite color or favorite movie), the school(s) from which they graduated, and the name of their first pet. Users were encouraged to fill out the questionnaire, send it back, and forward it to ten of their friends.



TIP Many websites use cognitive passwords created with multiple personal questions during registration. When users want to retrieve a forgotten password, the system queries them with these questions. As long as users answer the questions correctly, the system allows them to reset their password.

People who answered were at risk of giving out their answers to secret questions at banks and other online sites. For example, if you forget your password or log on to a bank's website from a different computer, the site will often challenge you. You're required to give the name of your first pet, for example, and if you can answer correctly, the system grants you access. If attackers have this information, they can use it to hack into your accounts.

Employees can accidentally provide inside information about their company when they spend time on these social networking sites. Because of the potential damage, many organizations prohibit the use of these types of sites on company computers. The rules are spelled out in acceptable use policies, which employees read and acknowledge at least once a year.

Information Rights Management

Information rights management (IRM) refers to the different methods used to protect sensitive information from unauthorized access. IRM is a subset of digital rights management (DRM). DRM is much broader and typically refers to protecting media such as music and video. IRM focuses on documents. It allows people to specify access permissions to different documents that they share, even after the documents leave their control.

For example, a user can protect a document with IRM before e-mailing it to an associate at a different company. IRM permissions can be applied to the document before it is sent, restricting what the recipient can do. IRM can prevent recipients from forwarding, copying, modifying, printing, faxing, or pasting the content if that's what the sender wants to restrict. It's also possible to restrict users from copying the content using the

Print Screen functionality and to set expiration dates on files. After the expiration date, users are unable to open the file.

Understanding Databases

Databases are stores of information that users can easily access and modify as needed. Large databases require sophisticated database management systems (DBMSs) to manage database access. Some common DBMSs in use today include products from Oracle, IBM, and Microsoft. These systems provide access to relational databases.

Understanding Database Tables

In a relational database, data is stored in tables, and this data is linked to other tables based on predefined relations. As an example, Figure 11-2 shows a Parts table and a Suppliers table with different elements identified. The Parts table holds information on parts, such as parts used for manufacturing a bicycle. The Suppliers table holds information on the company that supplies these parts. The SupplierID field links the two tables.



TIP Actual databases will have many more than just two tables. Many production databases used within organizations have hundreds of tables. A data dictionary describes the purpose of each table and its relationship with other tables. The ellipsis (...) in Figure 11-2 indicates that the table can have many more columns and many more rows, but for brevity, the figure only shows a few of the columns and rows. Some tables have dozens of columns, and they can have millions of rows.

Figure 11-2
Tables in
a database

Parts Table					
	Primary Key		Column or Attribute	Foreign Key	
	PartID	Part Name	QtyOnHand	SupplierID	...
	123	Widget	12	745	
Row or Tuple	124	Sprocket	23	745	
	125	Wheel	15	745	
	126	Horn	4	723	
	...				

Field

Suppliers Table			
	Primary Key		
	SupplierID	Supplier Name	Phone
	723	We Make Noise	555-765-4321
	745	Acme Supplies	555-123-4567
	...		

Key elements of relational databases include the following:

- **Row or tuple** Each row, or tuple, contains a unique data element. For example, the row for a sprocket identifies details about the sprocket, including how many sprockets are currently on hand in the QtyOnHand column.
- **Primary key** The primary key is used to uniquely identify each row or tuple. It ensures that each item has only a single row and that the item can easily be identified using the primary key. A primary key is normally an integer, but it can be any string of characters. In the Parts table, the primary key for the sprocket is 124, which will always reference the sprocket part and nothing else.
- **Column or attribute** Columns, or attributes, provide additional details on the data. For example, the Parts table could have additional columns, such as the cost of the part, a description, the color, and more.
- **Foreign key** A foreign key provides a relationship to another table. Notice that three of the parts have the same SupplierID of 745. The Suppliers table has the details of the supplier with a SupplierID of 745 (Acme Supplies). The foreign key in the Parts table points to the primary key in the Suppliers table. This allows you to enter the details for the suppliers once in the Suppliers table, instead of multiple times (in each row of the Parts table).
- **Field** A field provides an individual piece of information within any row or tuple.

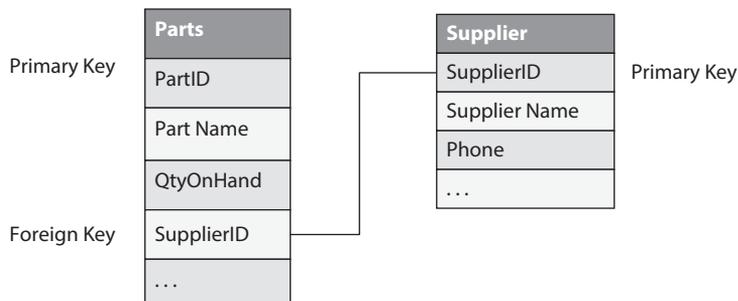
Understanding Relations

Tables are related to each other using the primary and foreign keys. The primary key ensures that each row or tuple within a table is unique. The foreign key in one table points to a unique primary key in another table to create the relationship. For example, Figure 11-3 shows the foreign key in the Parts table linked to, or related with, the primary key in the Suppliers table.

One of the benefits of the primary key and foreign key relationship is that data doesn't need to be repeated. For example, if you look back at Figure 11-2, you can see that the supplier with a SupplierID of 745 is supplying the first three parts. The Suppliers table shows that SupplierID 745 is Acme Supplies. If you entered all of the information on Acme Supplies (such as name, phone number, address, sales contact, and so on) in each row of the Parts table, it would become a lot of needless work. You would have to repeat

Figure 11-3

Table relationships



the same information in each of the rows, and each time that you reentered the data in a different row, the likelihood of making mistakes would increase.

However, by creating the relationship between the tables, you need to enter the data only once. If any of the data on Acme Supplies ever changes (such as a different sales contact or phone number), you only need to make the change once in the Suppliers table. If the supplier data were within the Parts table, you'd need to make the change in every relevant row or tuple of the Parts table.

Data normalization is the process of separating data into multiple related tables to reduce duplication of data. Duplicating the same data in more than one table creates dual entry (or multiple entry) problems, but data normalization eliminates these issues.



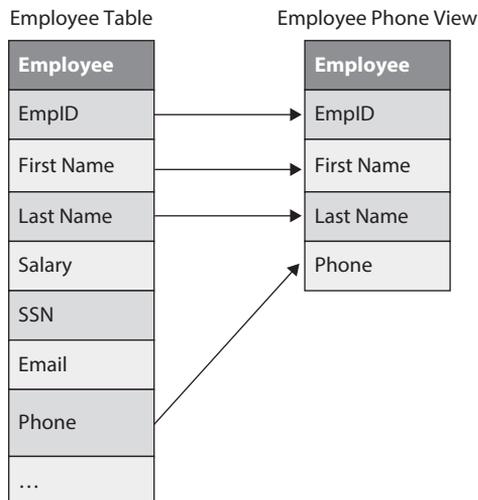
EXAM TIP Rows within a database are known as tuples. A primary key uniquely identifies each row in a table. A foreign key in one table points to a primary key in another table to link the tables together through a common relationship.

Using Views

One way that databases provide security is with views. A *view* is a virtual table that provides access to specific columns in one or more tables. It doesn't actually hold any data but presents the data in the underlying table or tables. A database administrator can grant access to a view without granting access to a table, thereby limiting what a user can see and manipulate.

For example, imagine that people in the company need to access phone numbers for employees and this information is contained within a table named Employee in a company database. However, the Employee table holds more information (such as salary data), and administrators don't want to grant users access to this additional information. Instead of granting access to the table, a database administrator can create a view and grant users access to the view, as shown in Figure 11-4.

Figure 11-4
Database view



Notice that the Employee table has other information such as salary and Social Security number (SSN) data that should remain confidential. The view shows several columns from the table, including the employee's name and phone number, without showing the confidential data.

Additionally, database administrators can grant access based on the users' needs. For example, an administrator can grant access for users to read or modify the data in the view. A key point is that users will only be able to read and modify columns in the view. When users modify data in a view, they modify the data in the underlying table. However, the view still restricts their access to other columns in the underlying table.

Communicating with Databases

Relational databases have their own language used to communicate with them, known as Structured Query Language (SQL). Although different companies often include their own unique additions to enhance SQL, they share some commonalities. SQL statements are primarily divided into two categories:

- **Data Definition Language (DDL)** DDL statements are used to create the structure of the database, including tables, primary keys, and relationships among the tables.
- **Data Manipulation Language (DML)** DML statements are used to add, modify, retrieve, and delete data from the databases. For example, a SELECT statement retrieves data, an INSERT statement adds data, an UPDATE statement modifies data, and a DELETE statement deletes data.

OLTP vs. OLAP

Databases are organized differently depending on how they're used. There are two primary methods:

- **Online transaction processing (OLTP)** Websites that access backend databases use OLTP. For example, each time a user makes a purchase, the web application records the transaction in the database as it occurs. OLTP databases need to be quick, and data normalization is one technique that ensures they provide the best performance. Additionally, OLTP databases record all transactions in logs before making the changes in the database. This allows the database management system to reverse failed transactions before writing them to the database.
- **Online analytical processing (OLAP)** Many organizations have very large databases known as data warehouses. Retrieving usable data from these data warehouses often requires additional steps. OLAP reorganizes data from data warehouses into multidimensional cubes for easier retrieval of useful information. These cubes are not normalized, but instead include redundant data. They are typically slower than OLTP databases, but speed isn't the most important consideration. However, business managers are able to access a significant amount of information and convert it into actionable data from OLAP databases. For example, a manager might want to compare how many blue widgets the

company sold last December to how many green widgets the company sold last December. This helps the manager determine how many widgets of each color to preorder for this December. This process is also known as *data mining*.

Data Inference

One of the security issues with data is *data inference*, the ability of someone to gain knowledge by piecing together unclassified data to determine classified or secret information.

As an example of data inference, the U.S. Air Force once had a problem with snipers taking potshots at jets as they flew sorties out of a certain airfield in a foreign country. The Air Force sent patrols out when sorties weren't planned, but the snipers only appeared when the jets were flying sorties. Even though the times and dates of the sorties were highly classified, the enemy learned of these sorties by accessing unclassified information.

Because the times and dates of the sorties were not consistent from day to day or even week to week, pilots and crew were notified when to report to work through schedules posted on a bulletin board. A foreign worker was able to view these bulletin boards and passed the information on to other foreigners working as spies collecting information. The spies were able to deduce that the posted work times corresponded to the times of the sorties. Although this is a lesson that is sometimes forgotten, military organizations recognize that any piece of information can be valuable to an enemy and take steps to protect all information. The U.S. Navy had a simple phrase to emphasize this during World War II, "loose lips sink ships," as part of a larger campaign known as "careless talk costs lives."



EXAM TIP Data inference attacks attempt to collect public or unclassified pieces of information to predict or guess an outcome. They can also attempt to gather large quantities of information and attempt to learn details of the information through deduction.

Data inference can sometimes work in the reverse also. In other words, if people have access to summary data (also called aggregate data), they can sometimes use this to learn details. For example, if personnel in accounting have access to total payroll data for any given day, they wouldn't necessarily know the actual pay for an individual. However, examining the total pay for the day before a person is hired with the total pay the day after a person is hired may allow someone to deduce the actual pay of the new hire.

Data Diddling

Data diddling is the unauthorized changing of data before entering it into a system or while entering it in a system. It is a form of fraud, and anyone involved with data entry has the opportunity to change it.

As an example, imagine Bob works in accounting and handles payroll. He is responsible for manually approving and entering overtime hours for employees. While entering the data, he adds more hours for his friend, which increases his friend's paycheck. Similarly, he and his friend could decide to do this together. Bob would increase his friend's overtime hours, and his friend would give him half of the extra pay.

Securing Big Data

Big Data refers to huge data sets that are so large that traditional methods of manipulating them are not adequate. Instead, developers often need to develop new methods to analyze the data. As an example, the National Energy Research Scientific Computing Center (NERSC) in Oakland, California, has a database with approximately 2.8 petabytes (PB) of data. Developers at NERSC have developed different tools to search, retrieve, store, and analyze this data. While 2.8PB is clearly large, it's worth noting that Big Data doesn't refer to a specific size.



TIP A petabyte is equal to 1 million gigabytes. The order is kilobyte, megabyte, gigabyte, terabyte, petabyte.

Personnel use the same processes to protect Big Data as they use to protect other databases. This includes using strong authentication and authorization processes for access control, and using encryption to protect the confidentiality of sensitive data.

However, because the applications that are used to access Big Data are new, they might have unknown vulnerabilities. The best way to discover these vulnerabilities is to test the applications. Additionally, because Big Data data sets are so large, they are frequently stored among multiple systems. This might require unique network architecture configurations that, being unique, might have unknown vulnerabilities. When developing applications and network architectures to support Big Data, it's best to use a process such as the system development life cycle (SDLC) described later in this chapter.

Regulatory Requirements

Laws mandate the protection of some data. This includes personally identifiable information (PII), health- or medical-related data, and other types of data. This section covers some of the regulatory requirements related to different types of data.

Personally Identifiable Information

PII is information that personally identifies an individual. The National Institute of Standards and Technology (NIST) Special Publication (SP) 800-122, *Guide to Protecting the Confidentiality of Personally Identifiable Information (PII)*, defines PII as

any information about an individual maintained by an agency, including:

- (1) any information that can be used to distinguish or trace an individual's identity, such as name, Social Security number, date and place of birth, mother's maiden name, or biometric records; and*
- (2) any other information that is linked or linkable to an individual, such as medical, educational, financial, and employment information.*

Although this definition is focused on a federal agency, it applies equally well to any organization that holds information on employees or customers that can identify an

individual personally. This also extends to information such as credit card data, any type of account access data, and driver's license numbers.

Several laws mandate the protection of PII, and if an attack compromises PII, organizations have a responsibility to inform the affected individuals. Some businesses consider these laws unnecessary and a barrier to business growth because they add expenses. On the other side of the fence, some people suggest that without laws forcing companies to disclose this information, they may not readily do so.



EXAM TIP Several laws mandate the protection of PII. Companies face serious fines if they do not protect PII. Additionally, if a data breach occurs, companies are obligated to let customers know of the breach.

For example, in April 2011, Sony suffered a data breach related to Sony PlayStation user accounts. The first result was a quick shutdown of the PlayStation Network, followed by several press releases starting about a week later.

This resulted in a significant amount of interest from government sources. The U.S. Energy and Commerce Committee's Commerce, Manufacturing, and Trade Subcommittee began investigating the attack, and the subcommittee's top members sent Sony a letter demanding answers about the incident. The U.S. Homeland Security Department was also involved in investigating the attack.

In one public statement to customers, Sony revealed, "We believe that an unauthorized person has obtained the following information that you provided: name, address (city, state, zip), country, e-mail address, birthdate, PlayStation Network/Qriocity password and login, and handle/PSN online ID." The company later released an e-mail that included this statement: "While there is no evidence that credit card data was taken at this time, we cannot rule out the possibility. If you have provided your credit card data through PlayStation Network or Qriocity, to be on the safe side we are advising that your credit card number (excluding security code) and expiration date may also have been obtained." This compromise affected as many as 77 million customers.

This example raises a question that people have debated: Did Sony reveal the information due to legislative requirements or out of genuine concern for its customers?

HIPAA

The Health Insurance Portability and Accountability Act (HIPAA) of 1996 mandates protection of health-related data within the United States. Prior to HIPAA becoming law, it was possible for just about anyone to call your doctor to ask for information on your health, and the doctor's office may have given it out freely. Today, any organization that gives out health-related data about a patient without authorization is subject to serious fines.

HIPAA identifies protected health information (PHI) as any information concerning the health status, provision of health care, or payment of health care for an individual. Although this primarily affects medical organizations such as hospitals, laboratories, and doctor's offices, it also extends to other, nonmedical organizations. For example, any organization that offers a health care plan and collects information about employees related to the plan also falls under the HIPAA rules.



EXAM TIP HIPAA covers any U.S. organization that processes health information. Failure to follow the requirements of HIPAA can result in significant fines and penalties.

SOX

The Sarbanes-Oxley (SOX) Act of 2002 mandates specific protections for data related to publicly held companies (companies registered with the U.S. Securities and Exchange Commission). SOX requires high-level officers (such as CEOs and CFOs) to verify personally the accuracy of financial data. Part of the goal is to prevent some of the fraud that was apparent in companies such as Enron and Tyco International.

For example, the Enron scandal resulted in a loss of \$11 billion to shareholders after the company's stock price plummeted from over \$90 per share to less than \$1 a share when it went bankrupt. Many experts attributed the failure to corrupt accounting practices. Later, several Enron executives were indicted and sentenced to prison.

Because the officers must personally verify the data, organizations are required to take extra precautions to ensure the accuracy and integrity of the data. Both internal and external audits are also required to ensure that these publicly held companies are in compliance.

Epsilon Data Breach

In 2011, attackers breached the databases of Epsilon, a marketing company, gaining access to customer names and e-mail addresses. At the time, Epsilon provided outsourced e-mail services for several major companies, such as JP Morgan Chase, Capital One Financial, US Bank, and others. Epsilon sends out as many as 40 billion e-mails annually to customers on behalf of these companies, although it indicated that only 2 percent of its database was actually compromised.

Legal requirements mandate notifying customers after data breaches such as this, and Epsilon notified its customers (the major companies) of the incident. These companies were then obligated to send out e-mails to their customers notifying them of the data breach. Within the span of a week, many users received e-mails from multiple companies informing them of a potential breach.

A company doesn't really want to tell its customers, "Sorry, but the bad guys have your data," and without laws in place requiring that companies do so, it's questionable whether all of these companies would have sent out these notifications. However, because laws required companies to send the e-mails, many of them took the opportunity to inform users of phishing attacks and identity theft. For example, an e-mail from Capital One included this paragraph:

It's always a good idea to ignore any e-mail that requests your confidential account or login information. And don't forget, if you get an e-mail you think is suspicious, don't click any of the links. Just send it to us at abuse@capitalone.com. Then delete it.

Training

Higher authority sometimes mandates specific awareness and training requirements related to data. For example, Federal Information Processing Standards Publication 200 (FIPS Pub 200), *Minimum Security Requirements for Federal Information and Information Systems*, published by NIST, specifies that organizations must meet at least the following minimum requirements for awareness and training:

- (i) ensure that managers and users of organizational information systems are made aware of the security risks associated with their activities and of the applicable laws, Executive Orders, directives, policies, standards, instructions, regulations, or procedures related to the security of organizational information systems; and
- (ii) ensure that organizational personnel are adequately trained to carry out their assigned information security-related duties and responsibilities.

This compliance starts with the organization defining data classifications. The U.S. government has done this with data classifications such as Top Secret, Secret, Confidential, and Unclassified. Next, the organization ensures that users are aware of the different types of data classifications. Training helps users understand the classifications, how the organization protects the data, and the users' responsibilities and requirements for protecting the data. Audits (internal and external) often use documentation of training to help verify that an organization is complying with basic requirements.

Managing Assets

Asset management refers to the steps used to manage important assets within an organization. It helps the organization maintain control of its assets. In this context, assets include hardware, software, and data.

On a larger scale, asset management is often associated with configuration and change management. In other words, an organization identifies all of its systems, but it also has documentation that identifies the current configuration of these systems. When a system is changed, the change management process ensures that the configuration documentation is kept up to date.



TIP Asset management systems help an organization track its assets. It's common for asset management systems to track hardware and software owned by an organization.

Hardware

Hardware management entails identifying relevant hardware that arrives at a company, periodically inventorying the hardware, and removing the hardware from the inventory when it is disposed of. This includes items such as desktop computers, laptops, servers, routers, and switches.

Not all hardware is added to an inventory. For example, a mouse is very inexpensive and the labor required to inventory all the mice in an organization is much more expensive

than the cost of replacing a few stolen or lost mice. Internal policies drive the decision to add hardware to the inventory. For example, a company policy may state that all hardware valued at more than \$100 should be added to the inventory and anything less than \$100 should not be added.

There are many different ways to manage a hardware inventory. A basic method is to record the model and serial number of each piece of hardware, place that information into a spreadsheet, and periodically inventory the hardware. Small organizations may use this method. Another inventory method labels systems with bar codes when the hardware arrives. The organization then periodically scans all the systems to verify that they haven't disappeared.

Some organizations use radio-frequency identification (RFID) tags. The RFID tags transmit information on each inventory item, and RFID receivers can read data from the RFID tags. This is much easier than using bar codes. An individual has to scan each system manually with a bar code system. However, with RFID tags, an individual can just walk around the room with an RFID receiver to complete an inventory.

When an organization wants to combine the system with configuration management, it often has enterprise tools that it periodically uses to inventory the systems electronically. For example, an enterprise application will periodically check systems over the network to ensure that they all have relevant patches applied. This same application can be used as part of the hardware inventory.

Software

Some asset management systems track software. This helps an organization track operating systems and software applications throughout their life cycle. It also helps administrators ensure that they are not violating licensing terms. For example, if an organization purchased 100 copies of an application, it needs to be certain that the number of installed copies of the application is 100 or less. Installing 200 copies of the application would be the same as software piracy.

Just as enterprise applications can check systems for hardware inventory and patch management, many can also inventory the systems and identify software installed on each system. For example, Microsoft systems support Windows Management Instrumentation (WMI). By sending WMI queries to devices over the network, administrators can retrieve a significant amount of information on the systems. This includes the hardware, operating system version, updates, and installed applications. Enterprise applications also use WMI queries to retrieve this information.

Data

Adding data to an asset management system is more difficult because it's so easy to make a copy of files. However, it is possible to add media to asset management systems. This includes external drives and backup tapes.

It's also important to follow procedures mentioned earlier in this chapter related to handling data. This includes ensuring the data is properly classified, marked, and labeled.

Certification and Accreditation

Certification and accreditation are two separate processes used to test, evaluate, and approve systems for specific purposes. The certification process includes several steps to evaluate, describe, and test a system. This includes all of the security controls that are used to mitigate risks to the system. Once a system is certified, an accrediting authority provides a formal declaration, which approves the system operation.



TIP Systems are often operating during the certification process. However, there may be specific restrictions on what the system can be used for until it is formally accredited.

The certification and accreditation process provides a level of checks and balances for systems. A single entity doesn't certify and accredit systems. Instead, one entity certifies them, and then another entity has the responsibility of verifying that the information is accurate and, if so, accrediting the system. A management authority normally completes the accreditation. When the authority accredits the system, it takes responsibility for the risks associated with the system and any losses associated with the system. Because the accrediting authority takes responsibility for the approval, this discourages simple "rubber stamp" approval. Instead, the accrediting authority verifies the authenticity of information provided by the certification process. Of course, because an organization expects the accrediting authority to examine the data closely, personnel take extra steps to ensure the data's accuracy before submitting the certification paperwork.



EXAM TIP Within the U.S. government, a Designated Approving Authority (DAA) provides official accreditation by approving a system for operation at a specific level of risk. It does not guarantee that a system is free of risk.

The U.S. government uses several publications from the NIST SP 800 series as standards for the certification process. For example, NIST SP 800-47, *Security Guide for Interconnecting Information Technology Systems*, provides information on the certification and accreditation of interconnected systems.

Certification, Accreditation, and Security Assessments

In addition to NIST SP 800 documents, FIPS Pub 200 mandates specific requirements related to certification, accreditation, and security assessments, as follows:

- (i) periodically assess the security controls in organizational information systems to determine if the controls are effective in their application;*
- (ii) develop and implement plans of action designed to correct deficiencies and reduce or eliminate vulnerabilities in organizational information systems;*
- (iii) authorize the operation of organizational information systems and any associated information system connections; and*

(iv) monitor information system security controls on an ongoing basis to ensure the continued effectiveness of the controls.



TIP Chapter 10 mentioned FIPS Pub 200 in the context of accounting and configuration management. Recall that it specifies minimum requirements that federal agencies must meet in several different IT security areas. Nongovernment organizations often adopt these same standards as best practices for IT security.

Common Criteria

The Common Criteria for Information Technology Security Evaluation (commonly called Common Criteria or simply CC) is a framework used to evaluate systems. It provides assurances that the specification, implementation, and evaluation of a system's security has gone through a rigorous and standardized process. The Common Criteria was created with the help of several government organizations, including those of the United States, Canada, United Kingdom, France, and Germany. It's recognized as an international standard and has superseded several other standards.

Evaluation Assurance Levels (EALs) provide a level of assurance for a product or a system. EAL has seven levels that indicate the level of quality assurance steps that personnel have taken to ensure the reliability of security features:

- **EAL1** Functionally tested. The threats to security are low, so a high level of security assurance is not required.
- **EAL2** Structurally tested.
- **EAL3** Methodically tested and checked.
- **EAL4** Methodically designed, tested, and reviewed. Many commercial operating systems achieve EAL4.
- **EAL5** Semiformally designed and tested.
- **EAL6** Semiformally verified design and tested.
- **EAL7** Formally verified design and tested. This provides the highest level of assurance. It's applicable in extremely high-risk situations, where the assets have exceptionally high value or when the impact of a threat can result in catastrophic losses.



EXAM TIP The Common Criteria uses EAL levels to determine the level of assurance of a system or product. EAL1 is the lowest level of assurance and EAL7 is the highest level. Many commercial operating systems achieve EAL4.

There's a trade-off in both cost and time to achieve higher levels of assurance. For example, the U.S. Government Accountability Office (GAO) indicates that it can take between 10 and 25 months to develop a system at EAL4, and the cost can range between \$150,000 and \$350,000.

TCSEC Superseded

You may run across documentation that talks about the Trusted Computer System Evaluation Criteria (TCSEC) and all the different components of it. U.S. government entities previously used TCSEC and it was widely known as the “orange book” because its cover was orange. The orange book included specific standards used to evaluate and classify systems in use to process and store classified information, or considered for use to process and store classified information.

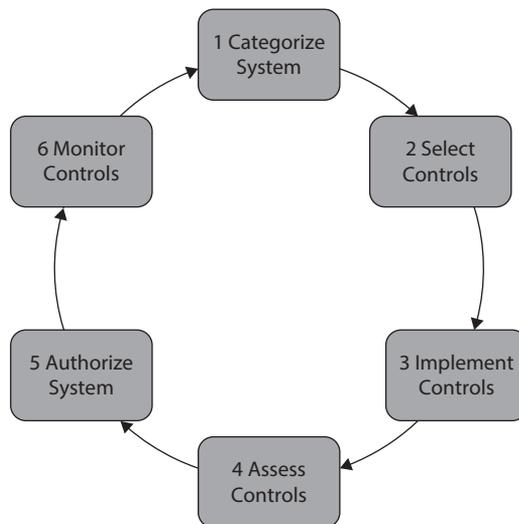
However, a Department of Defense directive in 2002 canceled the use of the orange book. It has been superseded by the Common Criteria since 2005. Studying TCSEC does provide interesting information on how standards have been developed over the years, but it isn’t relevant to certification and accreditation requirements currently used within the U.S. government.

Using a Risk Management Framework

NIST SP 800-37, *Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach*, provides guidelines on transforming the traditional certification and accreditation process into a six-step risk management framework. Figure 11-5 shows the steps in the framework, and the following text provides an overview of each of these steps:

1. **Categorize system.** Personnel examine the system, including the data it processes, stores, and transmits. They combine this information with an impact analysis to determine the security category of a system.

Figure 11-5
Risk management framework



2. **Select controls.** Personnel select a set of baseline security controls based on the category of the system. For example, a system processing confidential data would start with a baseline of controls to protect confidential data. Personnel then tailor or supplement these controls based on the system and the needs of the organization.
3. **Implement controls.** In this step, personnel implement the selected security controls. They also document information on the controls to show how they are used within the system and what risks they mitigate.
4. **Access controls.** Personnel periodically examine the security controls to verify they are implemented correctly, operating as intended, and working as expected.
5. **Authorize system.** If the level of risk is determined to be acceptable based on the assessment of the implemented controls, the system is authorized. Note that the system is often already operating before this formal authorization.
6. **Monitor controls.** Personnel monitor security controls on a continuing basis to determine their effectiveness against new threats and system changes. For example, any system change documented through a change control process may require a repeat of the six-step process.

The risk management framework recognizes that there are specific tasks required during different phases of the certification and accreditation process, but it also allows an ongoing reevaluation of the system on a regular basis. Risk doesn't stand still. Threats continue to morph, and mitigation controls implemented yesterday don't necessarily thwart the threats of today. Many security professionals believe that the risk management framework is better able to address these changes than the two-step certification and accreditation process.

Understanding Security Within the System Development Life Cycle

The system development life cycle (SDLC) is a model used to track systems (and sometimes software projects) from cradle to grave. In other words, the model identifies each of the different phases from the initiation of the project through to the ultimate disposition of the project at the end of its lifetime.

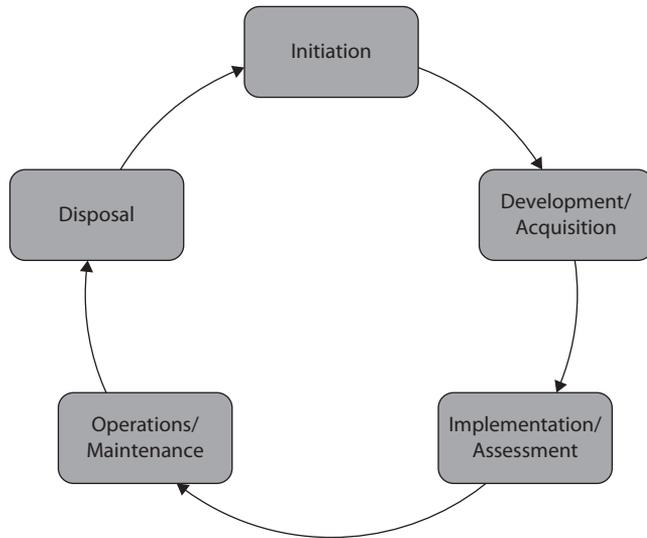


TIP Different sources sometimes show different phases for the SDLC. This discussion centers on the phases documented in NIST SP 800-64, *Security Considerations in the System Development Life Cycle*, of initiation, development/acquisition, implementation/assessment, operations/maintenance, and disposal.

Figure 11-6 shows the SDLC phases, and the following list provides a short overview:

- **Initiation** Personnel identify the need for the system in this phase. This includes documenting the purpose of the system and high-level requirements.

Figure 11-6
SDLC phases



- **Development/Acquisition** Personnel design, purchase, program, develop, or create the system. In many models, this phase includes other processes such as an acquisition cycle or a software development cycle.
- **Implementation/Assessment** Personnel install, configure, and test the system in this phase. Testing and evaluation steps determine whether the system meets the originally identified need.
- **Operations/Maintenance** During this phase, the system performs its primary mission. Personnel perform regular maintenance on the system to ensure it continues to operate as desired.
- **Disposal** In this phase, personnel remove the system from service. A key component of this phase is sanitizing all media. If it is being replaced by another system, it should be removed only after the other system has completed the implementation/assessment phase. Personnel should also update inventory records to reflect the final disposition.

While it's informative to understand the tasks required at different phases, much of this is outside the scope of the SSCP exam. From the perspective of the SSCP exam, the most important element is to understand the security issues at each of these stages. The following sections focus on the security elements.

Initiation Phase

Personnel focus on security concerns throughout the SDLC, and this includes the very beginning, the initiation phase. If security issues are not addressed early, then it's very possible that security fixes you try to add later will not be adequate or will cost too much

to implement. It's important to look at the data that will be processed, transmitted, or stored on the system and understand the value of the data to the organization. As a reminder, the system should be protected at least at the level of the classification of the data. By identifying the classification of the data that the system processes, the organization then knows the minimum level of protection required by the system.



EXAM TIP Security issues should be addressed starting in the initiation phase and continuing all the way through to the disposal phase. Ignoring security early in any project makes it much more difficult to address problems adequately later.

Development/Acquisition Phase

During the development/acquisition phase, an organization conducts a risk assessment to identify baseline controls to protect the system and its data. The risk assessment allows the organization to identify risks, develop security requirements to address these risks, and evaluate security controls to ensure they mitigate risks to an acceptable level.

Implementation/Assessment Phase

In the implementation/assessment phase, the organization installs the system and begins testing it. In addition to testing the system to ensure that it meets the needs of the organization, personnel test the security controls to ensure that they meet the required security specifications. If the security controls don't meet the requirements, it's appropriate to add more security controls. However, these new controls should go through the same testing process.

Operations/Maintenance Phase

The operations/maintenance phase is normally the longest phase of a system. One of the biggest risks during this phase is that changes to the system will reduce the effectiveness of security controls. Organizations use change management and configuration management practices to ensure that changes to a system do not reduce the system's security. Configuration management often includes an auditing component to check a system periodically. These checks not only ensure the system retains the required configuration, but also can detect unauthorized changes.

Disposal Phase

When the system's lifetime ends, personnel sanitize it prior to disposal. Also, it's common to migrate data from the system to ensure that the data will be available after the system has been disposed of.

Personnel use various methods to sanitize media, as described in the earlier "Removing Data Remnants" section. The method they use is dependent on the data held and processed by the systems. Higher classifications require more stringent sanitization methods, such as destroying the media. Personnel may use programs to remove files with lower classifications. The key here is that they need to sanitize all media to ensure that there are

not any data remnants on the media prior to disposal. This includes all disk drives, tapes, and nonvolatile memory. Systems should also be checked to ensure that optical media (such as CDs and DVDs) have been removed from the systems before disposing of them.



EXAM TIP If data remnants are not removed from systems, anyone who has access to the system after it has been thrown away will be able to read the data. Drives with data that has a higher level of classification may have to be destroyed, while drives with data of lower levels of classification may be scrubbed with programs that write different patterns of ones and zeros onto the media repeatedly.

During the disposal phase, it's also valuable to review the useful lifetime of data. Just because you can save data indefinitely doesn't mean you should. If data doesn't have any long-term use, it's often a good idea to dispose of it at some point. Of course, if requirements such as laws mandate keeping data for a specific period, it should not be disposed of until this time has passed.

Chapter Review

One of the first steps used to manage and protect data is to classify it based on its value to the organization. The classification drives the steps required to protect the data. Higher-classified data deserves greater protection, while lower-classified data does not warrant the extra costs associated with extra protection.

Data classifications in the U.S. government are Top Secret, Secret, Confidential, and Unclassified. Common classifications used in the private sector are Confidential (or Proprietary), Private, Sensitive, and Public, with Confidential representing the highest level. Many organizations implement data management policies to help protect data throughout its lifetime. Policies identify requirements when storing, transmitting, archiving, and retaining data. Access controls and encryption methods protect data at rest and data in motion.

Data is commonly stored in databases, and more specifically, in related tables within a database. Rows within a table are known as tuples. Tables have a primary key column to identify the row uniquely, and tables are related to each other by using a foreign key in one table to point to a primary key in another table. Database normalization separates data into multiple tables to prevent the duplication of data. SQL commands are used to communicate with databases. The two primary methods of organizing databases is for online transaction processing (OLTP) and online analytical processing (OLAP). Web e-commerce applications typically use OLTP databases because of their high transaction processing speed and store transactions in transaction logs. Managers use OLAP databases to retrieve actionable data.

Data inference is a risk with data that allows an individual to gather individual unclassified pieces of data to predict or guess an outcome. Similarly, inference may allow someone to view summary or aggregate data and identify details through deduction.

Personally identifiable information (PII) is protected through several laws, and organizations are often required to disclose whether control of any PII is lost or a database has been breached. The Health Insurance Portability and Accountability Act (HIPAA) mandates the protection of certain health information, and the Sarbanes-Oxley (SOX) Act mandates that executives of publicly held companies vouch for the integrity of a company's financial data.

Asset management ensures that an organization knows what hardware, software, and data it owns and helps protect those assets. It's common to tie asset management systems to configuration management and/or change management systems.

Certification and accreditation are processes used to test, evaluate, and approve systems for specific purposes. Certification includes steps to evaluate, describe, and test a system to identify and mitigate risks. After certification, an accrediting authority provides a formal declaration approving the system for operation. In the U.S. government, a Designated Approving Authority (DAA) provides the official accreditation of systems.

Some life cycle approaches to certification and accreditation are also being used. For example, NIST SP 800-37 provides a risk management framework used to identify specific tasks required during the certification and accreditation process. NIST SP 800-64 uses the system development life cycle model to track systems throughout their lifetime. When using any model, it's important to address security in every stage of the life cycle, including the very first stage.

Questions

1. Of the following choices, what type of data requires the least amount of protection?
 - A. Confidential
 - B. Public
 - C. Private
 - D. Sensitive
2. Which of the following provides the best confidentiality protection for data at rest?
 - A. Marking it
 - B. Labeling it
 - C. Backing it up
 - D. Encrypting it
3. Who is responsible for classifying data?
 - A. Management
 - B. User
 - C. Administrator
 - D. Owner

4. An organization wants to restrict risks associated with proprietary data transmitted over the network. What can it do in its data management policy to achieve this objective?
 - A. Restrict how long data is retained
 - B. Specify how data is deleted from storage media
 - C. Require the encryption of data in motion
 - D. Require the encryption of data at rest
5. Which of the following methods will reliably remove all data from a backup tape?
 - A. Erasing
 - B. Degaussing
 - C. Diddling
 - D. Sanitizing
6. Which of the following is a secure method of sanitizing optical media?
 - A. Degaussing
 - B. Overwriting
 - C. Shining
 - D. Destroying
7. A company wants to reduce the amount of space used to store files used and shared by employees. What can it use to reduce the amount of storage space used?
 - A. Data loss prevention (DLP) systems
 - B. Deduplication
 - C. Information rights management (IRM)
 - D. Retention policies
8. Users within an organization have recently sent sensitive data outside the organization in e-mail attachments. Management believes this was an accident, but they want to prevent a recurrence. Which of the following is the best method to do so?
 - A. Implement a network-based intrusion prevention system (IPS)
 - B. Provide training to users
 - C. Ensure the data is marked appropriately
 - D. Implement a network-based data loss prevention (DLP) system
9. Of the following choices, what is a tuple?
 - A. A column in a database
 - B. A row in a database
 - C. A primary key
 - D. A foreign key

10. Which of the following is a virtual table and allows a user access to a limited amount of data within a table?
 - A. View
 - B. Tuple
 - C. Row
 - D. Foreign key
11. An attacker has collected several pieces of unclassified information to deduce a conclusion. What is this called?
 - A. Data mining
 - B. Database normalization
 - C. OLAP
 - D. Data inference
12. An employee makes unauthorized changes to data as he is entering it. What is this?
 - A. Data diddling
 - B. Data entry
 - C. Data inference
 - D. Data deduplication
13. Information that can be used to distinguish or trace an individual's identity is also known as what?
 - A. PII
 - B. Tuple
 - C. Data inference
 - D. PHI
14. Which of the following choices identifies a regulation that mandates the protection of health-related information?
 - A. SOX
 - B. HIPAA
 - C. Epsilon
 - D. PII
15. The CEO of a publicly held company in the United States is required to verify the accuracy of a company's financial data. What requires this activity?
 - A. HIPAA
 - B. SOX
 - C. NIST SP 800-64
 - D. NIST SP 800-37

16. Within the U.S. government, who can formally approve a system for operation at a specific level of risk?
 - A. Certification authority
 - B. NIST
 - C. Senator
 - D. Designated Approving Authority (DAA)
17. Which of the following EAL levels indicates a system was methodically designed, tested, and reviewed, and is the level of assurance assigned to many commercial operating systems?
 - A. EAL0
 - B. EAL1
 - C. EAL4
 - D. EAL7
18. Which of the following is an international standard that provides a framework to evaluate the security of IT systems?
 - A. ITSEC
 - B. TCSEC
 - C. Common Criteria
 - D. Orange book
19. An organization is using a system development life cycle for the design of a system. When should personnel first address security issues?
 - A. During the initiation phase
 - B. During the development/acquisition phase
 - C. During the operations/maintenance phase
 - D. During the disposal phase
20. Of the following choices, what is a primary task to accomplish in the disposal phase of a system's life cycle?
 - A. Migrate all data to other systems
 - B. Delete all data
 - C. Remove data remnants from systems before disposal
 - D. Back up all data to tape

Answers

1. **B.** Public data requires the least amount of protection. An organization would want to ensure that public data on a website is not modified, but the organization puts the data on the website to make it available to the public. The other choices

indicate sensitive data that is important to an organization and deserves varying levels of protection.

2. **D.** Encryption provides the best confidentiality protection for data at rest. While it is appropriate to mark or label it, this isn't as strong as encrypting it. Backing it up provides protection for availability.
3. **D.** The data owner is responsible for classifying data. Management is responsible for defining data classifications, such as in a data policy. Users access the data, but they are not responsible for assigning or modifying data classifications. Administrators grant access to users based on the individual user's need and often at the direction from the data owner.
4. **C.** Encrypting data in motion can protect it against loss of confidentiality when it is transmitted over a network. Retention policies restrict how long data is retained, but do not affect data in motion. Destruction policies dictate how to delete data or destroy media, but aren't related to data in motion. Encrypting data at rest doesn't ensure that it is encrypted when it is transmitted. Some data-at-rest encryption methods will decrypt the data before transmitting it over a network.
5. **B.** Degaussing will reliably remove all data from a backup tape. Degaussing uses a powerful magnet to erase the data. Other methods of erasing data from a tape don't necessarily erase all the data. Data diddling is the unauthorized changing of data before or while entering it into a system and is unrelated to removing data from a tape. Sanitizing the tape is the goal, but it is not a method.
6. **D.** Optical media must be destroyed to ensure it doesn't include any remaining data. Degaussing isn't effective because optical media doesn't use magnetic methods of storing data. Overwriting isn't effective because the media might still have data remaining after overwriting it. Shining is not a valid method of sanitizing media.
7. **B.** Deduplication ensures that a file is stored only once on a system, even if multiple users have access to the same file. DLP systems attempt to monitor data usage and prevent the unauthorized use or transmission of sensitive data. IRM refers to the different methods used to protect sensitive information from unauthorized access. Retention policies restrict how long data is retained.
8. **D.** The best solution is to implement a network-based DLP system. It can scan outgoing data to look for sensitive data and block any it finds. An IPS focuses on incoming traffic to block attacks, and doesn't necessarily scan outgoing traffic. Thus, it wouldn't necessarily stop data sent as an e-mail attachment. Training is appropriate, but training doesn't necessarily prevent accidents. Also, training wouldn't stop a malicious insider from this action. The data should be marked with its appropriate classification, but there isn't any indication that it isn't. Also, even if the data is marked, it doesn't necessarily prevent accidents.
9. **B.** A tuple is a row in a database. Columns are also known as attributes. A primary key uniquely identifies a row in a table and is related to a foreign key in another table to create a relationship between two tables.

10. **A.** A view is a virtual table that provides access to specific columns in one or more tables, allowing a user to access a limited amount of data. A tuple is the same as a row and it contains a single record, but it is not a virtual table. Tables are linked together with foreign keys.
11. **D.** Data inference occurs when someone is able to piece together unclassified information to predict or guess an outcome. Data mining is the process of retrieving relevant data from a large database. Database normalization is the process of dividing the data into multiple tables in a database to reduce duplication of data. Online analytical processing (OLAP) uses techniques to make it easier to retrieve data.
12. **A.** Data diddling is the unauthorized changing of data before entering it into a system or while entering it into a system. While the employee is involved with data entry, data entry by itself doesn't indicate the employee is making unauthorized changes. Data inference occurs when someone is able to put together unclassified pieces of information to predict or guess an outcome. Data deduplication stores a file only once on a system, even if users attempt to store copies in multiple locations.
13. **A.** Personally identifiable information (PII) is any information that can be used to distinguish or trace an individual's identity and includes items such as their name, Social Security number, birthdate, birthplace, and more. A tuple is a row in a database table. Data inference occurs when someone is able to piece together unclassified information to predict or guess an outcome. Protected health information (PHI) is information referring to an individual's health and is protected by HIPAA.
14. **B.** The Health Insurance Portability and Accountability Act (HIPAA) mandates the protection of health-related information. The Sarbanes-Oxley (SOX) Act mandates specific requirements to guarantee the accuracy of data held by companies registered with the U.S. Securities and Exchange Commission. Epsilon is a marketing company that suffered a data breach and was required to report the loss to customers based on laws that protect personally identifiable information (PII).
15. **B.** The Sarbanes-Oxley (SOX) Act requires high-level officers (such as the CEO and CFO) to verify personally the accuracy of a company's financial data. HIPAA mandates the protection of protected health information. NIST publishes documents that provide guidance to federal agencies and can be used by private companies, but these documents do not require any actions by private companies.
16. **D.** A DAA provides official accreditation by approving a system for operation at a specific level of risk. The certification authority does not approve a system, but instead evaluates, describes, and tests a system. The National Institute of Standards and Technology (NIST) provides recommendations of standard best practices, but it does not certify or accredit systems. Senators do not certify or accredit systems.

17. **C.** Most commercial operating systems achieve Evaluation Assurance Level 4 (EAL4) when evaluated by the Common Criteria (CC). EAL4 indicates the operating system has been methodically designed, tested, and reviewed. EAL0 is not a valid level of CC. EAL1, the lowest level of assurance, indicates the system has been functionally tested. EAL7 is the highest level of assurance and indicates the system has a formally verified design and has been tested, but EAL7 ratings are not as common as EAL4, especially for operating systems.
18. **C.** The Common Criteria provides a framework for evaluating the security of IT systems, and several countries (including the United States, Canada, United Kingdom, and France) have adopted it. The Information Technology Security Evaluation Criteria (ITSEC) is a standard used in Europe, but it has been largely replaced by the Common Criteria. The Trusted Computer System Evaluation Criteria (TCSEC or orange book) was used by the U.S. government, but it has been superseded by the Common Criteria.
19. **A.** It's important to address security during each phase of the system development life cycle, starting with the initiation phase. If you start addressing security later, it's very possible that it will be more difficult and more expensive to add security controls.
20. **C.** Data remnants should be removed before disposal to ensure that data does not fall into the wrong hands. Depending on the value of the data, it could be migrated, deleted, or backed up, but this does not remove the need to ensure that data does not remain on the system before disposing of it.

This page intentionally left blank

Security Administration and Planning

In this chapter, you will learn about

- Security policy contents and characteristics
- Raising the awareness of security policies
- Business continuity plans
- Business impact analysis
- Disaster recovery plans
- Difference between a BCP and DRP
- Alternative locations such as hot sites, cold sites, and warm sites
- Security organizations such as NIST and US-CERT

Understanding Security Policies

A security policy is a written document that provides the organization with a high-level view of the security goals of an organization. Chapter 9 presented and contrasted the differences among policies, standards, guidelines, and procedures. To quickly review, policies are high-level, authoritative documents. Standards document criteria such as a proven norm or method and may specify requirements for a process or a technology, but the organization chooses what standards to follow. Guidelines provide recommendations for members of an organization, but they are not mandatory. Procedures provide the action steps to accomplish tasks.

A security policy is a management (or administrative) control that focuses on the management of risk and IT security. Senior management creates or at least endorses the organization's security policy. The policy reflects the culture of the organization and provides the overall direction for security. Others within the organization use the security policy as an authoritative document to implement controls to enforce and follow the requirements of the security policy.



NOTE A security policy can be an overall security policy for an organization. Organizations can also create security policies for individual systems that define the system and security protections related to the system. Within this chapter, the focus is on an overall security policy for an organization.

A security policy typically goes through several stages:

- **Initial stage** Personnel draft the security policy based on the needs of the organization. This might be a formal, nearly complete draft of the policy or an initial proposal to senior management identifying the needs and objectives of the policy.
- **Approval stage** Senior management approves the policy. It might take several iterations between the initial stage and the approval stage to create a document that senior management approves. Once approved, it provides direction for personnel to enforce the policy.
- **Publication stage** The policy is provided to relevant personnel so that they can follow and implement it.
- **Implementation stage** A security policy is an important first step to provide security within an organization, but it isn't the final step by any means. After senior management has approved a security policy, personnel must take steps to implement and enforce it.
- **Maintenance stage** Periodic reviews (such as once a year) ensure that the policy remains up to date, meets the needs of the organization, and addresses current threats.



EXAM TIP A security policy is a management (or administrative) control that provides the high-level authority for an organization. They are authoritative in nature and provide direction to personnel that implement them using other security controls.

While it's possible to have a single security policy covering all the elements of security for an organization, this can easily become very large and difficult to maintain. Instead, organizations often create several security policies, with a single security policy driving the others. As an example, Figure 12-1 shows the relationship between an organization's security policy and other supporting policies. For instance, the security policy can mandate the creation of an acceptable use policy (AUP) and require users to review and acknowledge the AUP periodically. Personnel then create an AUP to support this requirement. Similarly, the security policy may state that the organization needs to prepare for potential disasters and ensure that critical functions continue to operate during disasters. Personnel then create a business continuity plan, identify critical functions with a business impact analysis, and create disaster recovery plans to respond to different types of disasters.

Some organizations use a similar approach but include the supporting policies as appendixes to the security policy. The primary document is the organization's security policy, and it refers to the appendixes. Each appendix addresses separate security concerns, such as an acceptable use policy, an auditing policy, and so on.

Security Policy Characteristics

Depending on the size and needs of the organization, a security policy can be just a few pages long or 50 pages or more. There isn't a single standard for creating a security policy.

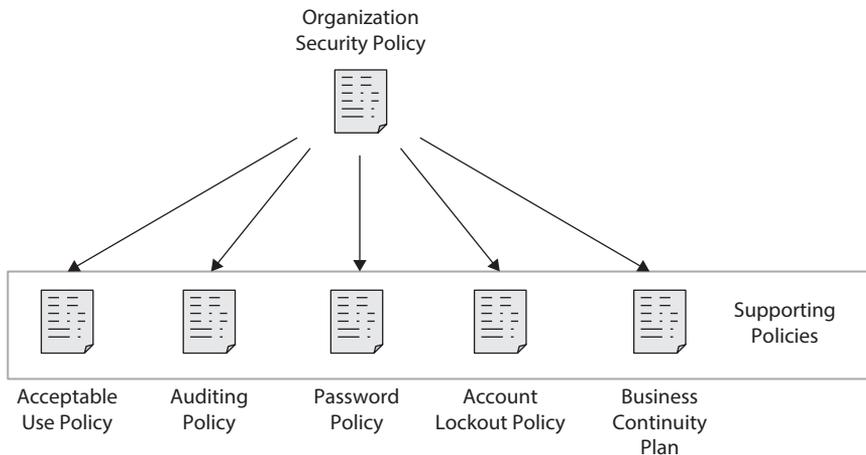


Figure 12-1 An organization security policy and supporting policies

However, you'll see some common characteristics in many effective security policies, including the following:



TIP Items in this list are not mandatory items to include within a security policy, and the list is not all-inclusive. However, the list does provide a cross-section of topics that many organizations include in their security policies.

- **Organization mission statement** A basic mission statement helps members of the organization understand the overall vision of the organization and how security can be used to enhance the mission and vision. The goal is for security methods and procedures to support the mission of the organization and its overall business model.
- **Statement of accountability** A security policy defines the roles and responsibilities of users, management, and IT staff. This section provides a basis for ensuring compliance with the policies and procedures.
- **Data classifications** An organization often defines the classifications of data (such as Confidential, Sensitive, Private, and Public) in the security policy. This section explains how to determine data classifications and provides some guidance on protecting data based on the classification.



NOTE Chapter 11 covers data classifications in greater depth. It includes a discussion of the common classifications used in the U.S. government and classifications commonly used in private organizations.

- **Backup plans** Security policies often dictate the requirements for overall backups. This includes backup requirements, offsite storage of backup copies, and how long to retain backups.
- **Classification of resources** A security policy can identify specific resources that need higher levels of protection. As an example, consider an organization that hosts an e-commerce website with access to a customer database. The security policy can mandate the protection of the e-commerce website and require the database to be stored on a different server. Administrators choose the details of how to implement these requirements. They might choose to host the web server within a demilitarized zone (DMZ) protected with an intrusion detection system, and store the database on a separate server in the internal network.
- **Network access** This section defines the types of access authorized on the network. For example, the organization may restrict the use of any wireless devices on a network. If the organization allows wireless access, the policy may specify minimum encryption standards. Similarly, the policy might mandate the creation of a separate network with Internet access for wireless devices, but prohibit these devices from accessing internal network resources.
- **Risk management** In this section, the organization defines an overall view of risk management. This could include direction on risk assessments and vulnerability assessments, such as how often to complete these assessments. An organization may also provide direction on types of analysis that assessors should use. For example, the policy might direct the use of quantitative analysis or identify specific values to use in a qualitative analysis.
- **Auditing** An auditing policy defines what should be audited at a minimum and how often. For example, it can mandate extra auditing requirements for privileged accounts and require that nonadministrator personnel audit actions by these accounts at least every six months.
- **Business continuity** Most organizations include an overall statement regarding business continuity and disaster preparation. Separate documents provide the details, but the security policy provides the overall goal.
- **Incident response** An incident response section defines a security incident for an organization and provides information on what should be done in response to an incident. It could include overall direction on goals to prevent, detect, and respond to security incidents and require the creation of an incident response team.
- **Physical security** This section provides an overview of physical security requirements. It may mandate the creation of a zone document for the building. The zone document identifies the specific protections required when processing different types of classified data and specific zones within the building. For example, a server room may process confidential data and require protections such as a single video-monitored entrance/exit that can be accessed only with a smart card and PIN.

- **Acceptable use** Most organizations draft acceptable use policies to let users know what is acceptable use of computer equipment and networks. The security policy may mandate the creation of an acceptable use policy and identify how users acknowledge it, such as by signing an acknowledgment that they have read the policy. The policy will often specify when or how often users should acknowledge it, such as when they are initially hired, and annually thereafter.



EXAM TIP An acceptable use policy (sometimes called an acceptable usage policy or AUP) ensures that users understand their responsibilities when using an organization's IT resources. Organizations commonly require users to review and acknowledge these types of policies regularly, such as once a year.

- **Enforcement section** Many policies identify an overall vision of how a security policy is to be enforced. For example, the policy may state an overall goal of using electronic monitoring and technical controls to enforce policy elements whenever possible. Similarly, it might mandate periodic training to ensure personnel are aware of the security policies.
- **Passwords and other authentication requirements** A policy can generically state the overall authentication requirements. For example, it may state that a technical password policy must ensure that users create strong, complex passwords and that they change them regularly, or that access to certain systems requires the use of two-factor authentication. Many organizations require administrators to have separate administrator accounts and use longer and more complex passwords (such as 15 characters or longer). When services require separate service accounts, these can have different password requirements, such as a minimum of 30 characters.
- **Account lockouts** The policy can require that accounts automatically lock after users enter the incorrect password too many times. Further, it can require technicians or administrators to unlock the accounts manually.
- **Hardware usage** The security policy can restrict or prohibit the use of some hardware. For example, USB flash drives have specific risks such as the ability to transfer malware from system to system, or allow an employee to steal massive quantities of data. Many organizations restrict the use of USB flash drives to reduce these risks. A security policy can mandate the use of technical methods to prevent the usage of USB drives altogether, or to prevent users from reading data from a USB drive or writing data to a USB drive. The policy can also require technical controls to detect and report any user attempts to use a USB drive.
- **Ethics statement** An ethics statement identifies the minimal acceptable behavior by members of the organization. The security policy can include an ethics statement identifying the organization's expectations related to ethics and directing an employee of what to do when faced with an ethical dilemma. In general, most organizations expect employees to resolve ethical dilemmas within the organization by going to their supervisor, or their supervisor's supervisor, whenever possible.

Enforcing Security Policies

Of course, security policies are useful only when personnel know about them, implement them, follow them, and enforce them. This is one of the reasons that senior management must endorse and support security policies. If not, the policies become meaningless pieces of paper.

Figure 12-2 shows some of the methods used to enforce security policies. You'll find these controls covered throughout this book. For example, Chapter 2 covers authentication and access controls; Chapter 7 covers risk management topics; and Chapter 9 covers several different types of management, technical, and operational controls.



TIP The elements shown in Figure 12-2 aren't meant to be a complete list. However, they do help illustrate that an organization uses multiple methods to support a security policy.

These elements work together to provide a defense in depth security solution by increasing the overall security posture of the organization. An important point worth stressing is that, ideally, the security policy comes first, and then personnel implement security controls in support of the security policy. The goal is for these controls to enforce the security policy and support the overall mission of the organization.

Value of a Security Policy

The obvious value of a well-crafted security policy is that it increases security and security awareness. If the security policy is able to prevent or at least reduce security incidents, it can protect an organization against direct financial losses, as well as indirect losses through damage to its reputation. However, there are also some less obvious benefits.

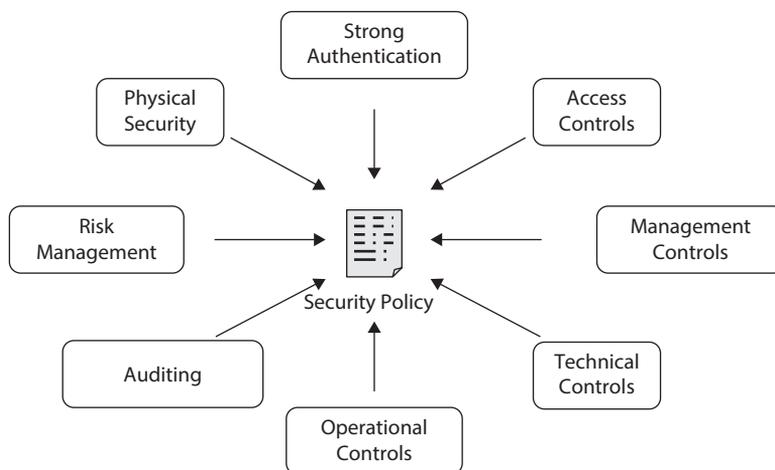


Figure 12-2 Enforcing a security policy

Direct Losses vs. Indirect Losses

Direct losses refer to the immediate losses resulting from an incident, such as the loss of revenue. Indirect losses refer to other factors such as the loss of customer goodwill.

As an example, attackers accessed personal information from approximately 110 million accounts during an attack on Target Corporation (a large U.S. retailer) in late 2013. Target later announced that the attack resulted in \$148 million in direct expenses, including money it set aside for existing and potential claims against the company.

Target also suffered additional, indirect losses. Profits fell 46 percent in the last quarter of 2013 (right after the attack) and its share price fell 4.4 percent. Following the attack, Target spent an additional \$61 million on security technologies and \$100 million more to implement chip-and-pin systems for more secure credit card transactions.

Even if a security incident occurs, the security policy may limit the organization's liability. As an example, consider a similar attack against two different fictitious organizations, ACME and EMCA. The attack exploits a zero day vulnerability and results in the loss of customer data. ACME doesn't have a security policy, and personnel have not documented their security practices. EMCA does have a security policy, and has a significant amount of documentation showing that it regularly follows security practices. ACME will be more vulnerable to litigation claiming that the company has not taken steps to protect customer data. In contrast, EMCA may be able to demonstrate that it did everything it could to prevent losses, so its liability should be limited.

Similarly, when employees sign security policies (such as an acceptable use policy) to acknowledge that they have read the policies and are aware of their contents, they can be held accountable to follow the security policies. The majority of employees won't knowingly violate security policies. However, some employees may purposely try to circumvent them, which can cause additional harm to an organization. For example, if an employee surfs pornographic sites using a company computer (counter to the AUP), fellow employees may see the objectionable material and raise sexual harassment issues. In the absence of a policy, the company may be liable for the sexual harassment. With a policy in place, the company can hold the offending employee accountable.

Security Policies Becoming More Common

Organizations are increasingly required to have security policies, typically because a higher authority mandates them. For example, the Payment Card Industry Data Security Standard (PCI DSS) requires organizations processing credit card data to meet multiple control objectives and security requirements, as detailed in Chapter 10. One of the requirements is to maintain a policy that addresses information security.

Some organizations aren't required by law or regulation to have a security policy and they consider it too expensive or too time-consuming to draft, implement, and enforce. This view often changes after one or two costly security incidents. The lack of a security policy can often cost much more in the long run.

Without a policy providing guidance, employees must use their own judgment to decide what they think is important to the organization. At the very least, these decisions will be inconsistent and can easily result in the loss of confidentiality, availability, and integrity of IT systems and data. Worse, the lack of direction often results in wasted efforts, rework, and higher costs.

Data is often one of the most valuable assets within an organization, but not all data has the same value. Without guidance, employees could easily overprotect public data and underprotect highly confidential data. For example, if IT personnel don't recognize the value of research and development data, they may not use strong access controls to protect it or may not back it up often enough. An attacker may be able to hack into the system and steal the data, or a common hardware failure may result in the loss of a significant amount of research data. Both scenarios can result in the organization's losing a significant amount of money.

It's also worth noting that if an organization doesn't have a security policy, it may be susceptible to other legal problems, including liability issues. For example, if an organization doesn't have an acceptable use policy and an employee uses a computer to hack into other systems or networks, the company can be held liable. However, an organization reduces its legal liability by having documentation (such as a signed AUP) proving it has informed employees of unacceptable practices.

Understanding the Code of Ethics

Chapter 1 covered the Code of Ethics required by (ISC)². As you'll recall, you must subscribe to this Code of Ethics to earn and keep the SSCP certification. The Code of Ethics includes responsibilities such as protecting society, acting honorably, providing diligent and competent service, and advancing and protecting the profession.

However, the (ISC)² doesn't have a corner on ethics. Many organizations have their own ethics statement that provides guidance to personnel related to their conduct within the organization. When employees sense a conflict with the organization's ethics, they have a responsibility to raise this internally within the organization rather than reporting it externally.



EXAM TIP Employees have an obligation to abide by the ethics of their employer. In other words, employees have an implied loyalty to their employer as long as they remain employed.

It's unlikely that any security policy will cover every possible issue that can present itself within IT security. This is where ethics come in. They provide a framework of norms and principles of correct conduct beyond the security policies. With this in mind, members of an organization have a responsibility to ensure that secure practices are still followed even if there isn't a specific policy that covers a particular challenge.

One point worth stressing is that employees should not retaliate against attackers, even if an organization doesn't have a policy addressing counterattacks. The exception is when it is part of the IT professional's job. For example, if an attacker launches a SYN flood attack on one of your servers, you may be able to analyze the logs and determine

the source of the attack based on the IP address. You may be tempted to launch an attack against the attacker, but this is not ethical. It can also result in unexpected results.

First, the attacker may be spoofing an IP address. Even though you think you know where the attack originated, you may be attacking an innocent company server that the attacker hijacked and used to launch the attack. Second, when you launch an attack, you are escalating the situation. If you block the attack, the attack stops and the incident is over. If you attack back, the attacker may take it personally and launch multiple malicious attacks against your organization for days, months, or even years afterward.

Policy Awareness

After an organization creates and approves a security policy, it's important to ensure that the appropriate personnel know of its contents. Having a policy hidden away in a desk is almost the same as having no policy at all.

Different members of an organization require different levels of understanding of a security policy. In other words, regular employees need to understand some topics within the policy, but not necessarily all the topics that security and IT personnel need to understand. As an example, all employees must be aware of an acceptable use policy.

In contrast, imagine the policy states that two-factor authentication is required when accessing the network via a virtual private network (VPN). IT and security personnel would design and implement an appropriate two-factor authentication system (such as with smart cards and usernames and passwords). Personnel using the VPN would need smart cards to authenticate via the VPN. However, users who never use VPNs don't need to know about the policy requirement. Similarly, users who do use VPNs don't need to understand the technical details of multifactor authentication, as long they know they must use a smart card with their username and password.

The following sections cover some common methods that help ensure personnel know the relevant details within the security policy.

Easy-to-Read Language

The policy should be written in language that is simple to understand. If the policy is too complex for the target audience, they won't understand it and very likely will not follow it.

As an example, if the policy states that passwords must have at least 30 bits of entropy, most people won't understand what that means. However, personnel will understand the policy if it states that passwords must have at least eight characters, must include at least one character from all four character types (uppercase letters, lowercase letters, numbers, and special characters), and must not include the user's name.

Warning Banners

One of the ways that you can let users know of relevant security policy contents is with warning banners. Administrators can create warning banners that appear each time a user logs on. Warning banners can define who is considered an authorized user and warn unauthorized users against accessing the system.

Some organizations go further and provide information to users on their expectation of privacy within the warning banner. For example, many organizations regularly monitor user activity. A warning banner can remind users that tools are monitoring their activity to detect improper or illicit use of computer systems and networks. The warning banner could also list penalties for noncompliance with the organization's acceptable use policy. These warning banners remind users of their responsibilities and help serve as a deterrent.

Training Sessions

Some elements of a security policy may require training to implement. Organizations typically target training to different groups based on the needs of individuals within the group. For example, an incident response team may require forensic analysis training to ensure they are able to respond to incidents, identify what happened, and limit the potential damage.

Regular users might only need annual training on the acceptable use policy and periodic training on emerging threats. For example, attackers continue to create different types of attacks such as phishing attacks. Periodically educating employees about current phishing attacks can help employees recognize malicious e-mails and avoid malware infections.

Many organizations provide this type of training by using online methods, such as via an internal website. However, even a simple e-mail can inform users of a current threat. For example, one phishing attack sent a ZIP file to users as an e-mail attachment, with the e-mail message indicating that FedEx was unable to deliver a product to them and instructing them to open the file for more information. However, the ZIP file included malware designed to infect their system. Sending an e-mail to users within the organization to inform them of the phishing attack not only helps them to avoid opening the file, but also helps them to recognize how phishing attacks morph over time.

Security Flyers and Posters

Some organizations use security flyers and posters to raise security awareness and remind people of the security policies. The goal is often to get the users to recognize that security is their responsibility too, not just something the company does.



EXAM TIP Some common methods used to increase user awareness of security policy elements that are relevant to them include training, warning banners, and posters. At the very least, relevant portions of the policy must be understandable and accessible to employees.

Updating Security Policies

Security policies are dynamic documents. An organization that creates a security policy but never updates it can't expect the policy to remain effective. Organizations change as time moves forward. Similarly, threats and vulnerabilities change. By periodically reviewing and updating a security policy, the organization keeps up with these changes.

It's common to schedule periodic reviews of a security policy as a maintenance stage. For example, many organizations schedule annual reviews. During this review, the organization can examine changes to determine whether the security policy warrants an update. Changes in the industry, business practices, laws, threats, and many other aspects may require changes in the security policy.



EXAM TIP Security policies should be updated regularly, such as once a year. If an organization's security policy is updated, supporting policies should also be reviewed to ensure that they still provide the necessary support.

Additionally, an organization will often reexamine a security policy after a security incident. For example, if a security breach results in the loss of valuable data to an organization, personnel should examine the security policy and supporting practices. They may find that the security policy is inadequate and needs to be updated. Or, they may find that the security policy itself is adequate but personnel didn't follow the policy. In that case, personnel might need additional training, or the policy might need to be more accessible to personnel who are responsible for implementing the policy.

Understanding BCPs and DRPs

Business continuity planning includes all the processes and procedures to prevent the loss of mission-critical services for an unacceptable length of time. A business continuity plan (BCP) includes multiple elements, such as a business impact analysis (BIA) and a disaster recovery plan (DRP). A BIA identifies critical business functions and processes, and a DRP helps an organization restore critical systems to full functionality as quickly and efficiently as possible after a disaster.

Several events can result in service disruptions, including any of the following:

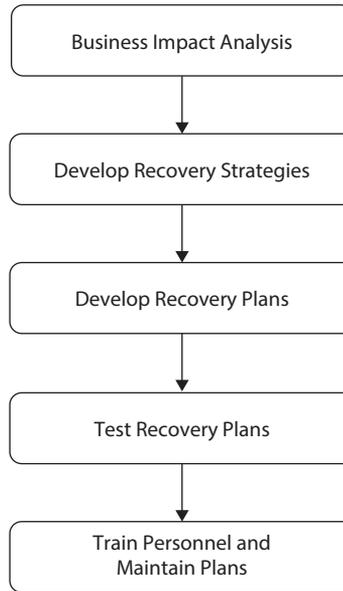
- Attacks
- Disasters
- Accidents
- Human error
- Equipment failure
- Software or application failure

It's worth stressing that business continuity planning isn't a small endeavor and it requires the support of senior management. This ensures that business continuity efforts are adequately funded and that all elements of the organization support the planning efforts. When organizations take business continuity planning seriously, it entails several distinct steps. Figure 12-3 shows the overall steps to develop a BCP.

The business impact analysis (discussed in more depth in the next section) identifies critical functions within the organization and provides an important starting point. The organization uses the results of the BIA to develop recovery strategies, such as the use of

Figure 12-3

Overall steps
for a BCP



alternative sites and different backup strategies. Once personnel identify recovery strategies, they develop and test plans to support the strategies. These plans include detailed steps used in response to disasters or emergencies. Last, personnel responsible for implementing the plans get appropriate training and begin implementing the plans. Personnel periodically review the plans and update them as necessary.

Figure 12-4 illustrates the relationship between a BCP and supporting policies. For example, a BCP starts with a BIA to identify critical functions. The BCP also includes backup policies outlining backup requirements and one or more DRPs that outline recovery strategies.

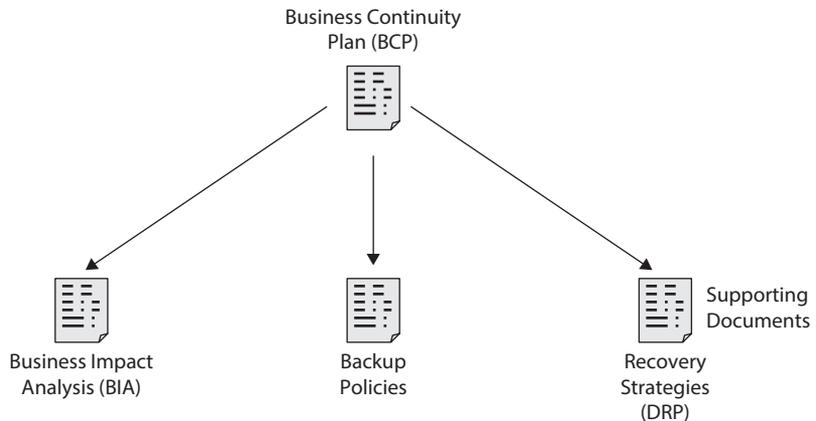


Figure 12-4 BCP and supporting elements



TIP The BCP is the primary document used for business continuity. It includes a BIA and other supporting documents such as backup policies and one or more DRPs.

Some organizations are in high-risk zones for specific types of events, so they often create supporting documents specifically for these events. For example, an organization that is in a high-risk zone for hurricanes would have a document specifically targeted at ensuring that critical operations continue during and after a hurricane.

Before the hurricane season begins, the organization dusts off last year's hurricane preparation plan and reviews it. It may decide that it wants to ensure that it is able to continue to operate critical functions for at least ten days after a hurricane hits. This drives other requirements, such as ensuring the location has at least ten days' worth of diesel fuel for generators and at least ten days' worth of food and water for personnel remaining onsite.

This plan would also identify steps to take as a hurricane is approaching, such as what to do when the hurricane is within 72 hours, 48 hours, 24 hours, and 12 hours of striking. For example, when the hurricane is within 72 hours, personnel check to ensure they have enough supplies on hand and begin bringing everything inside to prevent potential damage. When the hurricane is within 48 hours of the facility, the organization tests the generators and ensures that everything can continue to operate if the storm hits. When the hurricane is within 24 hours, emergency personnel report to work and all other personnel go home.



TIP The actual steps that the organization takes are identified in the hurricane preparation document and may be different from the steps identified in this timeline. However, the point is that the members don't have to guess about what they should do. Instead, they simply follow the checklists. Even if this is the first time that some of the personnel are experiencing a hurricane, they have the benefit of others' experience documented in these checklists.

Business Impact Analysis

A BIA is an important part of a BCP. It identifies the impact to the organization if any business functions are lost due to any type of incident. This information helps an organization identify what business functions are critical to continued operation. It also helps identify which resources are required to support these critical business functions.

For example, will the business stop functioning if a weather event causes a power outage in a building? The answer depends on what functions are performed in the building and whether any of these functions are critical. If this building houses primary servers used in an online e-commerce website, the loss of power could cause the entire business to cease functioning. In this case, the purchase of generators for long-term alternative power may be justified. However, if the building houses products and shipping facilities, a 24-hour or 48-hour power outage may result in only a minor impact on operations.

The BIA evaluates both direct costs and indirect costs. As mentioned previously, direct costs are associated with the immediate loss, such as an immediate loss to sales, while indirect costs are expenses related to recovering from the loss. For example, if an e-commerce website normally generates \$5,000 in revenue an hour, then the direct costs are \$5,000 for every hour the server is down. However, the organization may have to spend money on public relations to assure customers that their personal information is secure, and this extra money represents indirect costs. Additionally, some customers may leave forever, causing additional indirect costs.

Maximum Acceptable Outage

One of the most important findings of the BIA is the maximum acceptable outage (MAO) for different systems or services. This is sometimes called the maximum tolerable downtime (MTD) or the maximum tolerable outage (MTO). If an outage lasts longer than the MAO, it affects critical business functions supported by these systems or services.

As an example, consider an online website generating \$5,000 an hour during peak hours. After 60 minutes of downtime, it loses \$5,000 in direct revenue, so management may determine that it has an MAO of 60 minutes. This includes all the elements required to support the online website, including web servers, backend databases, Internet access, and the network infrastructure.

The same company might ship products from a separate shipping warehouse. Management may decide that warehouse outages as long as 48 hours will not affect critical business functions. As long as they communicate the delays with customers, the customers will accept some shipping delays. This results in an MAO of 48 hours for the shipping functions.



TIP Maximum allowable outage (MAO) is sometimes called maximum tolerable outage (MTO) or maximum tolerable downtime (MTD). All three terms refer to the maximum outage time before the outage affects critical business functions.

Recovery Time Objective

The MAO helps an organization determine the recovery time objective (RTO), which is the maximum amount of time that personnel can take to restore a system or service after an outage. If personnel take longer than the RTO to restore the system, the outage will affect critical business functions. If the MAO is 60 minutes, the RTO is also 60 minutes.

Administrative and security personnel also use the RTO to identify fault tolerance and redundancy methods. This includes Redundant Array of Independent Disks (RAID) subsystems and redundant servers configured in a failover cluster or load-balancing cluster. Implementing fault-tolerant and redundancy methods helps prevent the outages from occurring at all. For example, if a hard drive fails in a RAID, the system continues to operate. Similarly, if a server fails in a failover cluster, another server picks up the load and the service continues to operate.

These options can be expensive. However, if a web server generates \$5,000 in revenue an hour and the MAO is 60 minutes, management may decide that the cost is justified.

Recovery Point Objective

A recovery term often associated with databases is the recovery point objective (RPO). It identifies the maximum amount of data that an organization can afford to lose. This is sometimes difficult to identify. If you ask managers how much data they can afford to lose, you'll often hear "none" as the answer. However, it isn't until you identify the cost of restoring data up to the moment of failure that you get a more realistic answer.



NOTE Although an RPO is most often associated with databases, the same concepts can be applied to any type of backup.

If a database is relatively static, with just a few manual updates during the week, then management may determine that the RPO is a single week. In other words, backing up the database once a week is enough to meet this requirement. If a disaster deletes or corrupts the database, administrators can restore the database and personnel can manually enter the updates since the last backup.

In contrast, consider a database recording online transactions for a web server generating \$5,000 in sales an hour. Weekly backups aren't enough. If personnel back up the database on Sunday and a catastrophic database failure occurs on Wednesday, it results in the loss of all the customer transactions since Sunday. In this scenario, the RPO may be up to the moment of failure.

Although it's certainly possible to create a backup strategy for databases to recover a failed database up to the moment of failure, such a strategy can be expensive. That's why it's important to complete the BIA to help determine the actual needs of the organization related to the RPO. If the BIA determines the organization has a need to restore the database up to the moment of failure, it justifies the extra expense to support this RPO.

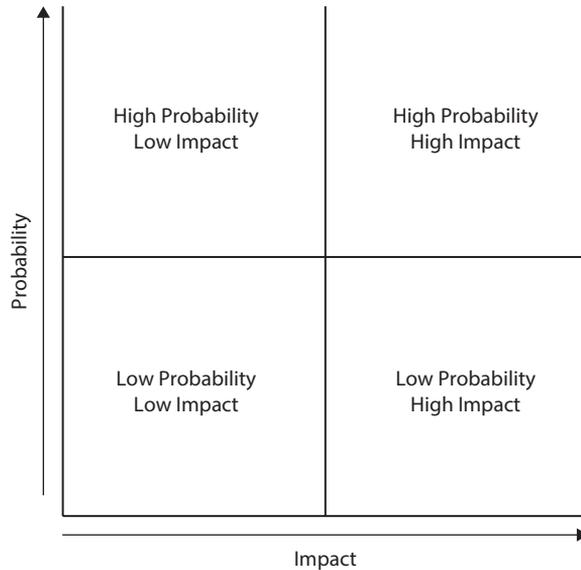


TIP Database backup strategies can include full and differential backups and the backup of database transaction logs. The transaction logs record every database transaction since the last backup (full or differential) and can be used to restore the database up to the moment of failure. Transaction logs are typically shipped to a different server for storage.

BIA Output

One of the major deliverables of the BIA is a document that identifies the actual monetary losses that can result from the outage of a critical business function. This is invaluable when you are evaluating controls that can mitigate these losses. You may remember from Chapter 9 that the controls are valuable only if they cost less than the losses they're trying to protect. For example, you wouldn't spend \$1,000 to protect a \$10 keyboard. Similarly, you don't want to spend business continuity funds protecting noncritical business functions. Although an e-commerce web server may be critical to an organization, an intranet web server used to track long-term projects may not be as critical.

Figure 12-5
Probability and
impact analysis



The BIA identifies how much money can be lost if the e-commerce web server fails. This provides management with an idea of the justified costs to keep the system operational even during a major incident. When evaluating multiple business functions, the BIA also attempts to prioritize them.

Figure 12-5 shows one way that a BIA can prioritize different functions by analyzing the probability of an incident occurring and the impact to the organization if the incident occurs. By graphing different types of disasters and identifying the systems that they can affect, an organization can identify the incidents that deserve the greatest attention.

Historical data helps an organization predict the probability of specific disasters. For example, an organization in San Francisco may consider the threat of an earthquake to be significant, while the threat of a hurricane may be nonexistent. Similarly, an organization in Miami may consider the threat of a hurricane to be significant, while the threat of an earthquake may be nonexistent. These are extreme examples, of course, and many other examples aren't so obvious. However, by analyzing and charting the results, the organization can prioritize different risks.

It's important to remember that historical data won't necessarily cover all potential negative incidents. Still, it is possible to predict some of these negative events. As an example, an organization might never have suffered any losses from a cyber-attack in the past. However, with the level of ongoing cyber-attacks, it's very likely the organization will be attacked at some point in the future. Personnel can predict the potential impact of a data breach based on the value of the organization's data.

Disaster Recovery Plan

A DRP provides an organization with a plan to restore critical operations after a disaster. The overall goal is to provide employees who are recovering the systems with clear-cut steps to restore the systems.

Disasters are major events, and people don't always think clearly when they occur. If employees have steps to follow within a checklist, they can ensure that they restore systems in a logical and orderly manner. On the other hand, if employees don't have steps to follow, they may spend a lot of time trying to determine what to do, and then more time redoing their steps when things don't work as expected.



EXAM TIP Disaster recovery and fault tolerance are not the same thing. Disaster recovery helps an organization recover after a disaster. Fault tolerance helps ensure that a system or component continues to function after a failure. Chapter 9 covered fault-tolerance strategies such as RAID and failover clusters.

Emergency Response Plans and Procedures

Disaster recovery plans include emergency response plans and procedures, and an organization will often have several different DRPs in place, ready to respond to different types of disasters. For example, organizations typically have plans for responding to disasters that can hit any organization, such as a fire or a crippling cyber-attack. Organizations will also have other plans that are based on local risks.

These additional plans address weather and nature-based events such as floods, tornadoes, hurricanes, cyclones, and earthquakes. While it's difficult to predict exactly when these events will occur, reviewing the history of an area can predict the likelihood that one of these events will occur at some point in the future. For example, no one should be surprised if a tornado touches down in Oklahoma during May. Similarly, no one should be surprised if a hurricane threatens land bordering the Gulf of Mexico during September or October.

Comparing a BCP and a DRP

It's important to recognize the differences between a BCP and a DRP. Many people combine them as though they are synonymous, but they aren't. In short, the BCP has a much wider scope and helps an organization continue to operate. In contrast, a DRP is a part of the BCP with a primary purpose of restoring a system to operation after an incident results in a major failure.

For example, if an organization is in the midst of a disaster and needs to move operations to an alternative location, the BCP would provide the details on the alternative location and what business functions it will perform. After the disaster occurs, the DRP would provide the details on how to restore operations for specific business functions.



EXAM TIP The BCP provides the information to keep critical functions running during a disaster (such as which critical functions to move to an alternative location). A DRP has a narrower focus and identifies how to recover one or more individual systems after a failure.

Even though a BCP and a DRP are not the same thing, you may run across documentation that treats them as though they are the same. However, (ISC)² documentation distinguishes between the two.

Restoration Planning

A key part of business continuity planning is restoration planning. Putting the concepts together, they typically occur in the following order:

- **Identify critical business functions** The BIA provides this information. Additionally, the BIA identifies the systems that support these critical business functions. For example, an online website might include one or more web servers, a backend database, Internet access, and a supporting network infrastructure.
- **Identify restore targets** The BIA identifies the MAO, and personnel use this to identify the RTO (and RPO if relevant).
- **Create plans to restore systems** These plans identify how to restore systems supporting critical business functions within the required times.

These plans need to be extremely clear and detail the exact order of restoration. For example, consider the online website with a backend database hosted on a database server. If the web application expects to connect with the database, it might have unexpected problems if it cannot connect to it. If that's the case, restoration plans might require technicians to restore the database server first, and then restore the web server. If the database is hosted on a group of servers within a failover cluster, the failover cluster might require servers to be restored in a specific order to ensure the failover cluster remains operational. Similarly, if the website is hosted within a load-balancing cluster of multiple servers, the cluster might require technicians to restore these servers in a specific order.

All of this might be common knowledge to the technicians and administrators who maintain these systems. However, you never know if these technicians and administrators will be available to restore the systems during a disaster. When the restoration plans provide clear directions on what to restore, in what order to restore it, and how to restore it, it makes it easier for any technician or administrator to complete the restoration in the shortest amount of time.

Testing and Drills

The first draft of any plan is rarely perfect. However, testing helps identify the problems so that they can be resolved before an actual disaster. NIST SP 800-84, *Guide to Test, Training, and Exercise Programs for IT Plans and Capabilities*, provides some excellent guidance on testing methodologies. It discusses tabletop exercises and functional exercises in-depth.

- **Tabletop exercises** These are discussion-based exercises where personnel meet in a conference room or classroom setting. A facilitator presents a scenario such as a hurricane or tornado and asks the exercise participants questions related to

their roles. Talking through the scenario gives participants a better chance to identify potential flaws in the written plans.

- **Functional exercise** Personnel actually perform the steps outlined in the plan as a method of validating the plan. Functional exercises can be focused on specific elements of a plan or be very broad in their scope. For example, you can give a technician a plan to restore a database server, and see if the technician is able to restore it on a test server using the plan. The actual database server remains online, so the test does not affect actual operations. The exercise can also encompass a full-scale exercise designed to address all the plan elements. While it is only a simulation, it allows personnel to perform their roles as they would in an actual emergency.

In some cases, plans require personnel to take specific actions. A classic example is a fire. If a fire occurs, personnel should exit the building using a predefined path, which helps the maximum number of people exit the building in the shortest amount of time. A simple way to test this type of plan is to hold a drill. The drill requires personnel to respond as if the emergency is occurring. In the case of a fire drill, people hear the fire alarm and exit the building. Fire departments are often happy to participate in these types of drills and provide feedback to the organization on how to improve their plans.

Alternative Locations

If an organization determines that it must be able to continue operations even if a location suffers a catastrophic failure, it designates an alternative location. For example, banks and investment companies need to provide access to customer data even if a weather event (such as a tornado, flood, or hurricane) causes so much damage to one location that they can no longer operate at that location.

Alternative locations include hot sites, cold sites, warm sites, and mobile sites. Figure 12-6 provides an overall view of each of these alternative location types, and the following sections explain them in more depth.

Hot Site

A hot site includes all of the resources necessary to take over operations of another location in a very short time, typically within an hour and sometimes sooner. A hot site includes hardware such as servers and the network infrastructure, up-to-date data, and personnel to manage the functions of the alternative location.



TIP An extreme version of a hot site is a mirrored site. A mirrored site is identical to the primary location and provides 100 percent availability. As data changes at the primary location, real-time data transfers send the updates to the mirrored site. This ensures that the mirrored site is always up and operational. In contrast, a hot site might take as long as an hour to become operational.

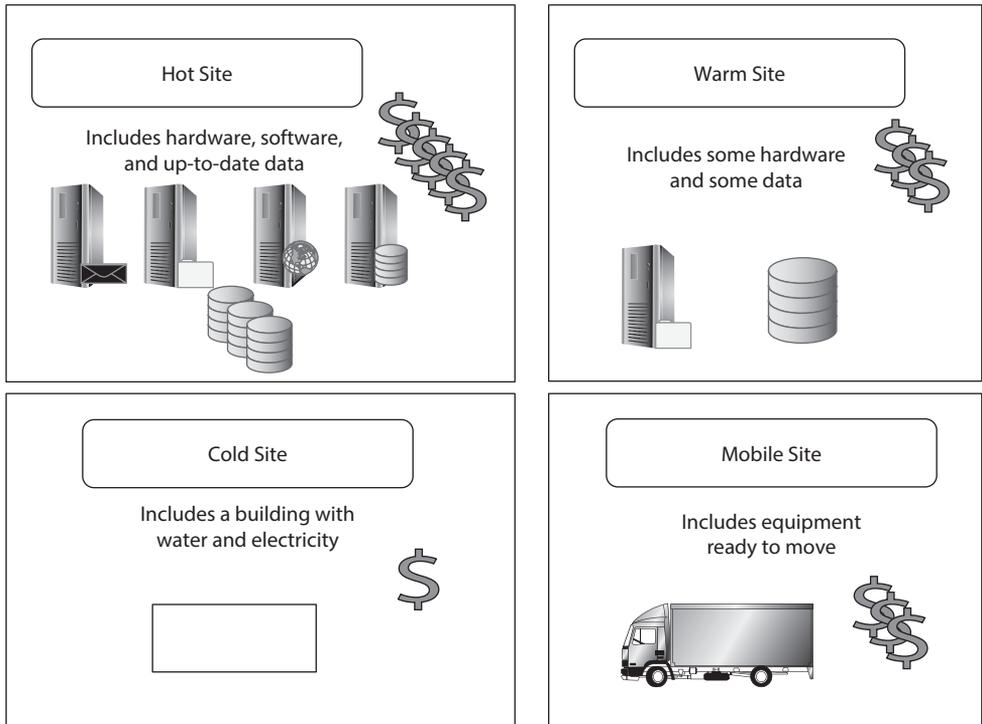


Figure 12-6 Comparison of alternative locations

Hot sites are valuable if an organization needs to prepare for a catastrophic failure that can occur at any time. For example, many banks and investment companies have hot sites ready to take over in case of major incidents, such as the terrorist attacks of September 11, 2001. Organizations that operate in earthquake or tornado zones often need to have hot sites available. Hot sites are the most expensive of the alternative location types, but they are the easiest to test, because everything is in place.

It's rare for a hot site to be used only as an alternative location for an organization. Instead, the organization typically uses it for less critical operations the majority of the time, but keeps it in a ready state to take over critical operations at a moment's notice. For example, the organization could use the location to provide basic customer services, but also host the equipment and infrastructure as a hot site. If an emergency requires the organization to activate the hot site, personnel at the alternative location scale back customer services and dedicate more resources to take over for the primary location.

Similarly, regional locations often have the capability of taking over another region's services. For example, an organization can have major locations in the western and eastern parts of a country. If either region's location fails, the other location can take over the critical operations of the failed location.

Cold Site

A cold site is a building with a roof, running water, and electricity. It doesn't include the necessary hardware, software, or personnel. In the event of an emergency, personnel move all of the resources to the cold site location, hook it up, and configure the site for operation.

It's difficult to test a cold site because nothing exists at the location. However, a cold site is significantly cheaper to maintain than other types of sites.

Warm Site

A warm site is a compromise between a cold site and a hot site. The organization makes compromises with costs and time. Instead of ensuring that the site can be activated within minutes, the organization may decide that 24 hours' notice is enough to bring it online.

For example, an organization may operate in a hurricane zone. Because hurricanes travel relatively slowly, communities often have more than a day's notice of a possible hurricane landfall. When a storm is approaching, the organization can use the time to activate a warm site. This includes deploying personnel to the warm site and updating all of the data.

Mobile Site

When an organization doesn't want to use a dedicated location as an alternative location, it can create a mobile site. For example, it's possible to set up the inside of a storage container as an alternative location. These are the same storage containers that 18-wheeler tractor-trailers use to haul goods, and they are as easy to move as any other loads hauled over the highways.

Organizations outfit these storage containers with all the equipment needed to support critical operations. If an emergency requires activating the mobile site, the organization moves it to the desired location and hooks it up to power. Depending on its needs, the organization might use an external heating and ventilation system, or bring one with the container and then mount it on the roof or outside the storage container. Similarly, the organization might connect directly to a broadband Internet service provider (ISP) for Internet access at the mobile location, or use a satellite connection.

Mobile sites like this will have all the necessary equipment, but the software and data are typically out of date. When the mobile site reaches its destination, personnel will ensure that the equipment is up to date with all appropriate patches and approved changes and that the data is up to date.

Mobile sites don't have to be as elaborate as a fully equipped storage container. For example, one small company in Virginia Beach has a simpler mobile site plan. When a hurricane is approaching, the owner packs the company's primary server into her car and drives west. When she settles somewhere, she plugs a USB wireless adapter into the server, configures her phone as a personal hotspot, and uses it to provide Internet access for the server.

The cost of a mobile site is similar to the cost of a warm site—somewhere in the middle of a cold site and a hot site. The actual cost varies depending on what an organization chooses to use for its mobile site. Clearly, the packed storage container will be much more expensive than the owner's car, smartphone, and a single server.

Identifying Security Organizations

Several security organizations exist that provide different levels of assistance and documentation for IT security personnel. It's worthwhile knowing about some of these organizations and their available documentation.

NIST

Several chapters within this book mentioned the National Institute of Standards and Technology (NIST). NIST includes a Computer Security Division and an Information Technology Laboratory (ITL). The organization regularly performs research in the area of information technology in direct support of federal agencies. However, NIST uses federal funds to perform this research, so it makes this research publicly available as Special Publications (SPs).

The SP 800 series of documents includes over 100 well-researched documents on a wide range of IT security topics, including cloud computing, risk management, intrusion detection systems, protection of wireless transmissions, protection of personally identifiable information (PII), and much more. You can view a full list of these documents (and download any of them) at <http://csrc.nist.gov/publications/PubsSPs.html>.



TIP NIST also publishes Federal Information Processing Standards (FIPS) documents, which define specific minimum requirements for many federal agencies. These documents often provide guidance for specific federal agencies, but they also provide sound guidelines for private companies.

Some of the documents that directly relate to security policies and continuity planning include

- **SP 800-18** Guide for Developing Security Plans for Federal Information Systems
- **SP 800-30** Guide for Conducting Risk Assessments
- **SP 800-34** Contingency Planning Guide for Federal Information Systems
- **SP 800-37** Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach
- **SP 800-41** Guidelines on Firewalls and Firewall Policy
- **SP 800-53** Security and Privacy Controls for Federal Information Systems and Organizations
- **SP 800-58** Security Considerations for Voice Over IP Systems
- **SP 800-61** Computer Security Incident Handling Guide
- **SP 800-64** Security Considerations in the System Development Life Cycle
- **SP 800-77** Guide to IPsec VPNS
- **SP 800-83** Guide to Malware Incident Prevention and Handling for Desktops and Laptops

- **SP 800-84** Guide to Test, Training, and Exercise Programs for IT Plans and Capabilities
- **SP 800-94** Guide to Intrusion Detection and Prevention Systems (IDPS)
- **SP 800-97** Establishing Wireless Robust Security Networks: A Guide to IEEE 802.11i
- **SP 800-100** Information Security Handbook: A Guide for Managers
- **SP 800-111** Guide to Storage Encryption Technologies for End User Devices
- **SP 800-113** Guide to SSL VPNs
- **SP 800-115** Technical Guide to Information Security Testing and Assessment
- **SP 800-122** Guide to Protecting the Confidentiality of Personally Identifiable Information (PII)
- **SP 800-123** Guide to General Server Security
- **SP 800-124** Guidelines for Managing the Security of Mobile Devices in the Enterprise
- **SP 800-127** Guide to Securing WiMAX Wireless Communications

NIST also publishes monthly bulletins available here: <http://csrc.nist.gov/publications/PubsITLSB.html>. Each monthly bulletin focuses on a single topic of interest related to computer security. For example, the March 2015 bulletin, titled “Guidance for Secure Authorization of Mobile Applications in the Corporate Environment,” outlines a process an organization can use to review and approve or reject applications used on mobile devices.

ITL vs. ITIL

The Information Technology Laboratory (ITL) is not related to ITIL. ITIL (formally known as the Information Technology Infrastructure Library) is a group of books documenting best practices for IT management. The United Kingdom Office of Government Commerce created the original ITIL books, but ITIL has gone through several changes over the years. Axelos, Ltd. currently owns all of the ITIL intellectual property, and it licenses the use of ITIL materials to other organizations. Axelos also accredits organizations as licensed Examination Institutes for the ITIL-related certifications.

In contrast, a U.S. government entity creates the ITL documents, which are freely available to anyone. The SP 800 series of documents are public-domain documents rather than intellectual property. Additionally, ITL doesn't sponsor any certifications.

US-CERT

The United States Computer Emergency Readiness Team (US-CERT) provides response support and defense against cyber-attacks for several government entities in the United States and for some industry and international partners. It is operated as part of the

National Cyber Security Division (NCSD) of the Department of Homeland Security as a public/private partnership.



TIP As many as 250 organizations use “CERT” in their name and deal with IT security incidents and responses. US-CERT is not affiliated with any of these organizations, although it does work with some of them. For example, US-CERT coordinates with the CERT Division of the Software Engineering Institute (SEI), a federally funded research and development center at Carnegie Mellon University.

US-CERT sponsors the National Cyber Awareness System, which includes several methods of providing information to users and technical experts. Chapter 6 mentioned several of the newsletters that US-CERT publishes, and you can sign up for any of them here: <https://www.us-cert.gov/ncas>.

SANS Institute

While the SANS Institute is a private institution that sponsors other security certifications and sells training for these certifications, it also publishes some free resources. For example, the SANS Reading Room includes more than 2,000 free white papers on a wide variety of computer security topics. You can access the SANS Reading Room here: www.sans.org/reading-room/.



NOTE SANS looks like an acronym but it is the name of the company. However, the name is derived from Sysadmin, Audit, Networking, and Security.

SANS also sponsors the Internet Storm Center, mentioned in Chapter 6. The Internet Storm Center monitors threats related to malware and provides free information via diaries (<https://isc.sans.edu/diaryarchive.html>) and podcasts (<https://isc.sans.edu/dailypodcast.xml>).

CERT Division

The first CERT Coordination Center was created at Carnegie Mellon University in response to the Morris worm that struck in 1998, and was later renamed as CERT Division. It is a federally funded program located in the Software Engineering Institute (www.cert.org/). The CERT Division works closely with the Department of Homeland Security (DHS) and helped the U.S. government create US-CERT. Both organizations share information with each other and work together on some issues.

Chapter Review

A security policy is a written document, authoritative in nature, that provides a high-level view of the security goals for an organization. External standards may influence the security policy, but the organization decides what standards to follow. Personnel develop procedures based on the guidance from the security policy.

Senior-level management should create, or at least endorse, the organization's security policy. Once the security policy is created, personnel within the organization implement and enforce it. Users must be made aware of the contents of the security policy through training, warning banners, posters, and other awareness methods, and they should be asked to read and sign an acceptable use policy. It's common to review a security policy periodically, such as once a year or after a security incident.

Business continuity plans (BCPs) include processes and procedures to prevent the loss of mission-critical services. A BCP starts with a business impact analysis (BIA) and includes one or more disaster recovery plans (DRPs). It's important to recognize that a BCP is not the same as a DRP, but a DRP is often a component of a BCP.

The BIA identifies the maximum acceptable outage (MAO) time (also called maximum tolerable downtime, or MTD) for critical services and systems. If an outage for one of the critical services or systems exceeds this time, it affects critical business functions. The MAO drives the recovery time objective (RTO), or the maximum amount of time that the organization can take to restore a critical service or system. The recovery point objective (RPO) refers to the point in time to which a database needs to be recovered if the system fails.

A DRP provides the steps required to restore a system after an outage. Organizations test both BCPs and DRPs using various testing methods. Tabletop exercises allow personnel to talk through the steps of a given scenario. Functional exercises simulate an event and allow personnel to walk through the steps of a given scenario.

Some BCPs require the use of alternative locations. A hot site includes all the hardware, software, and up-to-date data required to take over operations at a moment's notice. A cold site includes a roof, water, and electricity, but little else. A warm site is a compromise between a hot site and a cold site. Some organizations use a mobile site instead of a fixed location, which allows the organization to move its operations to another location during an emergency.

NIST includes a Computer Security Division and an Information Technology Laboratory. The ITL has published over 100 Special Publications in the SP 800 series on a wide range of IT security topics. All of these special publications are available for free download here: <http://csrc.nist.gov/publications/PubsSPs.html>. Other organizations that provide security-related materials include US-CERT and the SANS Institute.

Questions

1. Of the following choices, which provides the highest-level authority for an organization?
 - A. Standards
 - B. Policies
 - C. Guidelines
 - D. Procedures

2. Which of the following choices best describes an organization's security policy?
 - A. An authoritative written document that identifies an organization's overall security goals
 - B. A nonauthoritative written document that identifies an organization's overall security goals
 - C. A technical control that mitigates risks
 - D. A baseline used to ensure that systems are secure when deployed
3. Which of the following best describes the purpose of a security policy?
 - A. Ensures personnel understand their responsibilities
 - B. Ensures personnel use strong authentication
 - C. Informs personnel of management priorities related to security
 - D. Provides guidance on management controls
4. An organization wants to ensure that users are aware of their responsibilities related to the use of IT systems. What should the organization create?
 - A. A video monitoring system
 - B. An audio monitoring system
 - C. An acceptable use policy
 - D. An account lockout policy
5. Sally notices that Homer appears to be stealing from the company. What should Sally do?
 - A. Confront Homer
 - B. Ignore the activity because it doesn't concern her
 - C. Call the police
 - D. Report the activity to a manager
6. Which of the following choices are effective methods of ensuring that employees know the relevant contents of an organization's security policy? (Choose all that apply.)
 - A. Providing training
 - B. Using warning banners
 - C. Using posters
 - D. Storing the policy in the company vault
7. An organization has a security policy in place. What can personnel within the organization do to ensure it remains relevant?
 - A. Perform audits
 - B. Perform training
 - C. Review it
 - D. Test it

8. Which of the following is the most important element of business continuity planning?
 - A. Support from senior management
 - B. Availability of a warm site
 - C. The backup plan
 - D. Cost
9. What is the purpose of a BIA?
 - A. To identify recovery plans
 - B. To drive the creation of the BCP
 - C. To test recovery plans
 - D. To identify critical business functions
10. Which one of the following is a valid step to perform during a business impact analysis?
 - A. Identify alternative locations
 - B. Create a plan to restore critical operations
 - C. Identify resources needed by critical business functions
 - D. Identify minimum outage times for key business services
11. What is MTO in relation to business continuity planning?
 - A. Minimum time for an outage
 - B. Maximum time for an outage
 - C. Minimum tolerable outage
 - D. Maximum tolerable outage
12. Which of the following best describes maximum tolerable downtime?
 - A. The maximum amount of downtime before a business loses viability
 - B. The point in time in which a failed database should be restored
 - C. The maximum amount of time that can be taken to restore a system or process
 - D. The minimum amount of time that can be taken to restore a system or process
13. What is RPO in relation to business continuity planning?
 - A. Restoring potential outage
 - B. Recovery point objective
 - C. Restoration process option
 - D. Recovery process options

14. What is RTO in relation to business continuity planning?
 - A. Recovery terminal objective
 - B. Recovery time objective
 - C. Recovery tolerable outage
 - D. Recovery tolerable objective
15. An organization's location has been hit by a tornado and the organization is moving to an alternative location. What provides the direction for this action?
 - A. BIA
 - B. BCP
 - C. DRP
 - D. Hot site
16. An organization decides to designate an alternative location to be used in case of an emergency. The organization doesn't need anything other than an open building with water and electricity. What type of site best meets this need?
 - A. Hot
 - B. Warm
 - C. Cold
 - D. Distant
17. An organization decides to designate an alternative location for operations during a disaster. The site must be up and operational within minutes of an outage at the primary location. What type of site best meets this need?
 - A. Mobile
 - B. Cold
 - C. Warm
 - D. Hot
18. An organization is updating its business continuity plan (BCP) and wants to implement an alternative location that is the easiest to relocate. What type of site best meets this need?
 - A. Cold
 - B. Hot
 - C. Mobile
 - D. Warm

19. Of the following choices, what is a U.S. government entity that regularly publishes Special Publications (known as SP 800 series documents) related to IT security?
- A. ITIL
 - B. NIST
 - C. CERT Division
 - D. US-CERT
20. Which of the following organizations provides regular cyber-security alerts about current security issues, vulnerabilities, and exploits as part of the U.S. National Cyber Awareness System?
- A. ITL
 - B. NIST
 - C. CERT Division
 - D. US-CERT

Answers

1. **B.** Policies are high-level documents and provide authoritative direction. Standards can influence a policy, but an organization chooses what standards to follow. Guidelines provide recommendations, but they are not mandatory. Personnel create procedures based on the policy.
2. **A.** A security policy is an authoritative (not nonauthoritative) written document that identifies an organization's overall security goals. It is a management (or administrative) control, not a technical control or a baseline; however, technical controls and baselines are created based on the direction from the security policy.
3. **C.** Of the choices, the best description is that the security policy informs personnel of management priorities related to security. The other answers provide some specific goals of a security policy but do not address the overall purpose. An acceptable use policy helps ensure personnel understand their responsibilities. Technical controls ensure personnel use strong authentication, but a security policy covers more than just technical controls. Similarly, a security policy provides guidance on more than just management controls.
4. **C.** An acceptable use policy (also called an acceptable usage policy) lets users know what is acceptable use of computer equipment and networks. Monitoring systems ensure users follow policies, but they don't ensure users know the policy. Account lockout policies lock out users after too many failed password attempts.
5. **D.** An employee's loyalty should be to the organization, so Sally should report this activity to a manager or supervisor. An organization's ethics policy often includes procedures for reporting these types of incidents. Confronting Homer may not solve the problem but instead may result in Homer causing problems

for Sally, especially if Homer is stealing from the company. Because continued employment is based on the success of an organization, losses are the concern of every employee. The organization should make the decision of whether or not to call the police.

6. **A, B, and C.** Providing training, using warning banners, and using posters are all effective methods of ensuring that employees know the relevant contents of a security policy. If the security policy is stored in the company vault, it won't be accessible to employees.
7. **C.** A security policy should be reviewed on a regular basis (such as once a year or after a security incident) to ensure that it is still relevant. Audits help to prove that the security policy is being used and enforced. Training ensures that people know the contents of the security policy. It's appropriate to test a BCP or a DRP, but not a security policy.
8. **A.** Of the available answers, the most important element is support from senior management. While an organization might decide it needs a warm site, not all BCPs require warm sites. A security policy may mandate the creation of backup plans, but this is separate from the BCP. The cost is a concern, but the requirements drive the cost, and without support from senior management, business continuity planning may not receive any funding.
9. **D.** The business impact analysis (BIA) identifies critical business functions and is a part of the BCP. Personnel create recovery plans later in the process, after creating recovery strategies. The BCP drives the creation of the BIA, not the other way around as suggested by answer B. You can only test the plans after personnel have created them.
10. **C.** A core goal of a BIA is to identify critical business functions and the resources needed by these critical business functions. Identifying alternative locations is part of a business continuity plan (BCP). A disaster recovery plan (DRP) is a plan to restore critical operations. A BIA does identify maximum acceptable outage times, but not minimum outage times.
11. **D.** The maximum tolerable outage (MTO), sometimes called maximum allowable outage (MAO) or maximum tolerable downtime (MTD), identifies the maximum amount of time that a system can be down before critical business functions are affected. The *T* does not represent time, and the *M* does not represent minimum.
12. **A.** The maximum allowable outage (MAO), sometimes called maximum tolerable downtime (MTD), indicates the maximum amount of downtime a business can tolerate and still maintain viability as a business. Recovery point objective (RPO) indicates the point in time to which a failed database should be restored. Recovery time objective (RTO) represents the maximum amount of time that can be taken to restore a system or process after an outage. MTD is not related to minimum timeframes.

13. **B.** RPO represents recovery point objective and indicates the point in time to which a failed database should be restored. The other answers are not valid terms for RPO within business continuity planning.
14. **B.** RTO is an acronym for recovery time objective and represents the maximum amount of time that can be taken to restore a system or process. The RTO is derived from the maximum allowable outage (MAO). The other answers are not valid terms for RTO within business continuity planning.
15. **B.** The business continuity plan (BCP) provides direction for moving to an alternative location after a disaster; the primary purpose is to continue to provide critical business functions. The business impact analysis (BIA) helps an organization identify what functions are critical. The disaster recovery plan (DRP) has a narrower focus and helps an organization recover one or more systems after the disaster has passed. A hot site is a possible type of alternative location, so it doesn't provide direction for the action.
16. **C.** A cold site is a building with water and electricity and nothing else (such as no computer equipment or data). A hot site includes everything and is ready to take over operations within a short time after an outage. A warm site is a compromise between the two. A distant site isn't a valid term for an alternative location.
17. **D.** A hot site includes everything and is ready to take over operations within minutes after an outage. A mirrored site (not available as an answer) is a type of hot site that can take over operations almost immediately. A mobile site will take much longer to relocate and set up. A cold site is a building with water and electricity, and it would take much longer to become operational. A warm site is a compromise between a hot site and a cold site and takes longer to become operational than a hot site.
18. **C.** A mobile site is the easiest to relocate. Hot, cold, and warm sites use a designated location that the organization either purchases or leases.
19. **B.** The National Institute of Standards and Technology (NIST) publishes documents in the SP 800 series related to IT security. ITIL is a United Kingdom project. CERT Division is a federally funded program located in the Software Engineering Institute at Carnegie Mellon University, and the United States Computer Emergency Readiness Team (US-CERT) provides response support and defense against cyber-attacks. While CERT Division and US-CERT publish documents related to IT security, they don't publish SP 800 series documents.
20. **D.** The U.S. Computer Emergency Response Team (US-CERT) manages the National Cyber Awareness System, which provides cyber-security alerts, bulletins, tips, and updates. The other organizations provide information, but not alerts through the National Cyber Awareness System.

This page intentionally left blank

Legal Issues

In this chapter you will learn about

- Computer forensics
- Incident handling processes
- Three phases of a computer forensic investigation
- The importance of preserving and protecting evidence
- How to protect volatile RAM and data on disk drives
- The difference between computer abuse and computer crime
- How mandatory vacations and job rotation help prevent fraud
- Some of the laws and regulations designed to protect PII

Exploring Computer Forensics

Computer forensics is the science of examining and inspecting computer systems for evidence about an event or crime. The examination includes recovering data stored on or transmitted by any electronic device. A primary goal when identifying and recovering digital evidence is to preserve its integrity. This ensures that forensic experts can analyze it and that the evidence can be used in legal proceedings if necessary.

You've probably watched one of the CSI television shows where detectives use crime scene investigation (CSI) techniques to piece together the elements of a crime. Investigators are careful to ensure they don't modify the evidence as they collect it, and they protect the evidence after collecting it. Computer forensic experts have the same goals as they try to piece together evidence to investigate computer abuse or computer crime. The only difference is that computer forensic experts typically use different tools and techniques when they are investigating an incident. The following sections describe some of the important concepts to consider related to incident handling and computer forensics.

Participating in Incident Handling

As an IT or security professional, you are sure to observe security incidents, so you should be prepared to participate if asked. With that in mind, it's important to understand the overall incident handling process. Chapter 7 included a section on responding to incidents and covered methods documented in National Institute of Standards and Technology (NIST) Special Publication (SP) 800-61, *Computer Security Incident Handling*

Guide. This section covers the topics from a slightly different perspective using the terminology in the SSCP objectives.

As a reminder from Chapter 7, a security incident is any violation of policies or security practices that has the potential to result in an adverse event. Negative consequences include system crashes, unauthorized access to data, and the execution of malicious code on a computing system. Security incidents also include any violation of a security policy or an organization's approved security practices.

Many documents refer to the first step an organization takes in incident handling as preparation. Preparation includes implementing security controls (or countermeasures) to prevent, detect, and correct security incidents, as discussed in Chapter 9. Many organizations also create an incident response policy that outlines its plans to respond to an incident.

Once an incident occurs, the primary goal of incident handling is to contain any potential damage and then repair the damage. The following sections discuss the individual steps in an incident handling process.



TIP Incident handling processes are generally the same. However, different documentation labels these processes with different terms. Similarly, different organizations may label the processes differently in their security policies. Within an organization, the correct terminology is what the managers of the organization have documented within these policies.

Discovering an Incident

IT and security personnel often don't know an event is an incident until they respond and begin investigating. For example, a user may report erratic behavior on a computer. After investigating it, an IT administrator could discover malware infecting the user's system. In this case, the user's report is the event. Once the administrator discovers the malware, it has become an incident.

Responding to an Incident

After discovering an incident, personnel respond in accordance with the organization's incident response plan, which might include actions to take after discovering specific incidents. For example, it's important to contain an incident to limit potential damage. If an administrator discovers a computer is infected with malware, the incident response plan might specifically state that the IT administrator should disconnect the network interface card (NIC) to isolate the computer from the network. If the malware hasn't spread yet, this action contains the malware to the infected computer.

The incident response plan obviously cannot include specific actions for every possible incident. However, it will include generic guidance such as protecting evidence while investigating the incident, containing the incident, escalating the incident, and reporting the incident.

Escalating an Incident

Personnel escalate an incident if they realize they cannot resolve it. This is similar to a help desk technician escalating a help desk ticket when the technician cannot solve the

customer's problem. A more experienced technician takes over and attempts to resolve the problem. Similarly, if technicians or administrators discover a security incident but realize they cannot resolve it, they take steps to escalate it. Often, this simply means they report the incident to their supervisor or a higher-level administrator to get some help.

Personnel do not escalate all incidents. For example, if a technician discovers malware and uses antivirus software to remove it, there is no need to involve other personnel. However, the technician would still document it using the organization's procedures.

Reporting an Incident

Reporting an incident refers to informing personnel within the organization and entities outside the organization about the incident. Internally, an organization's incident response plan might require personnel to report any security incident to a member of the incident response team. As an example, technicians that identify and remove malware from systems might still be required to report the incident to security personnel, even though they resolved the incident.

Many laws and regulations require organizations to report certain types of incidents. This is especially true for any data breaches that expose personally identifiable information (PII) or protected health information (PHI). These laws typically require the organization to report the data breach to affected individuals. This helps ensure that people are informed when a data breach has exposed their personal information. They can take appropriate countermeasures such as subscribing to a credit monitoring service.

Implementing Feedback Loops

Feedback loops help ensure that an organization examines incidents in an attempt to learn from them. As a simple example, after discovering and resolving a malware infection, security personnel will often try to determine the steps leading up to the malware infection. They may find that the malware infection occurred for one of the following reasons:

- Antivirus software was not installed or was disabled, or signatures were not up to date.
- Up-to-date antivirus software did not block the malware.
- A user brought the malware in from home with a USB drive.
- A user opened an attachment in a phishing e-mail, which installed the malware.
- A user clicked a link in a phishing e-mail, taking the user to a malicious website.
- A user visited a malicious website, which downloaded malware in a drive-by download.
- A system didn't have up-to-date patches.

By understanding what allowed the malware infection, security personnel are better prepared to recommend countermeasures to prevent a recurrence. For example, the

following list identifies possible countermeasures security personnel can recommend for the malware infection:

- Use technical security controls to install and update antivirus software
- Evaluate the effectiveness of the antivirus software and consider replacing it
- Implement antivirus software at other locations, such as on e-mail servers and at the network boundary (between the Internet and the internal network)
- Use technical security controls to block access to USB drives except for those purchased by the organization
- Purchase antivirus software for users to install on their home computers
- Provide training to users on the dangers of transferring data between home and work computers via USB drives
- Provide training to users on phishing e-mail tactics
- Provide training to users on drive-by downloads
- Terminate employment of users that ignore the training
- Implement or improve a patch management program

Implementing Countermeasures

Management has the responsibility of selecting which countermeasures to implement based on the organization's acceptable level of risk. Once management approves the countermeasures, other personnel begin implementing them. Of course, it's possible that the organization will experience a similar incident even after implementing new countermeasures. This repeats the process. The result might be that the organization implements new countermeasures or replaces existing countermeasures.

First Responders and Preserving the Scene

In the medical field, the first responder is the first medically trained person to arrive on the scene of an incident that requires medical knowledge. It can include emergency medical services (EMS) personnel, emergency medical technicians (EMTs), police, firefighters, or other personnel. In the IT field, the first responder is the first forensics-trained person on the scene. This may be someone with advanced forensics training. It can also be an IT administrator who understands the basics of preserving forensic evidence and can manage the scene until a forensic expert arrives.

Medical first responders aren't expected to have the same level of knowledge as a doctor. Instead, they're expected to have enough knowledge to stabilize an injured person's health until a more qualified professional arrives or they are able to transport the injured person. Similarly, IT first responders aren't expected to have the same level of knowledge as full-time forensic experts. However, they should know the basics of handling evidence, preserving the scene, and maintaining a chain of custody.



EXAM TIP First responders are the first forensics-trained personnel on the scene of a computer incident. One of their primary responsibilities is the preservation of evidence. NIST SP 800-86, *Guide to Integrating Forensic Techniques into Incident Response*, discusses first responders in more depth.

Preserving the evidence is extremely important in a computer forensic investigation. If the evidence is not preserved, other personnel cannot verify the analysis. Additionally, evidence that isn't preserved is considered tainted. Tainted evidence can be misleading, indicating that innocent people are guilty and that guilty people are innocent. Tainted evidence loses its value, and a court of law won't allow its use in legal proceedings.

Along these lines, many organizations recommend disconnecting computers from a network as an initial step in a computer investigation. This isolates the system and prevents remote access. This ensures someone cannot clear logs or delete data on the system via a remote connection. It's also an important step to take if personnel suspect the system is infected with malware, such as a worm, that may infect other systems on the network.

Anti-Analysis Malware Destroys Hard Drives

Rombertik is a type of anti-analysis malware that checks for specific attributes associated with malware analysis. It checks the system's memory to determine if tools are analyzing it and, if so, attempts to destroy the master boot record on the system's hard drive. If this fails, it encrypts all the files on the C: drive.

Attackers typically deliver Rombertik via phishing e-mails that trick users into opening a malicious attachment. It includes several other methods designed to avoid detection. The primary purpose of Rombertik is to collect information from user systems and send it back to the attackers.

It's common for malware developers to use techniques to avoid detection. However, destroying the hard drive after detecting analysis is relatively new.

The best protection against this type of malware is a basic security control—reliable backups. With reliable backups, this type of attack is inconvenient, but it isn't catastrophic.

Three Phases of a Computer Forensics Investigation

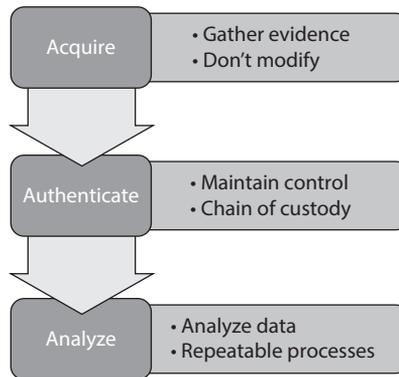
A computer forensics investigation is typically described as occurring in three phases of acquiring, authenticating, and analyzing the evidence, as shown in Figure 13-1. IT professionals need to be aware of the three phases and understand their role if an incident occurs. For example, if IT professionals sound the alarm to report an incident, they need to know what their responsibilities are related to protecting the evidence prior to security professionals arriving to secure the evidence.



EXAM TIP The ultimate goal of a computer forensics specialist is to acquire, authenticate, and analyze electronic evidence. It's extremely important that no one modifies the evidence during this process.

Figure 13-1

Three phases of a computer forensic investigation



Acquire the Evidence

In the acquisition phase, the investigators collect the evidence. It's very easy to modify the evidence during this phase if personnel don't take steps to preserve the evidence. As a simple example, opening a file modifies it because the file now shows a new access time.

Forensics experts use bit-copy or drive-imaging tools to capture the contents of a drive without modifying the data. This copy is a forensic duplicate image. In addition to copying data, this process includes copying free and slack space on the drive or partition. The copy must be an exact bit-by-bit copy of the original in order for legal entities to consider it forensically sound. For example, Figure 13-2 shows the overall process of copying a drive. Experts analyze the copy instead of the original drive. If they want to validate their analysis, they make another copy of the original drive and repeat their analysis. Additionally, the original drive remains available as evidence in a court of law if required.

In addition to capturing the contents of a drive, it's often useful to capture the contents of a computer's memory. This often has information on recently run processes and applications and the information can be useful during the investigation. However, the computer's memory is volatile and will be lost when the computer is turned off, so it's important to keep the system on until forensic experts have an opportunity to capture the memory contents. Forensic experts have other tools available that can read and copy data from volatile memory.



EXAM TIP It's important that anyone who responds to an incident does not access any files on a system or power a system down. These steps can modify or destroy evidence that may be valuable for an investigation. Before examining hard drives, you should copy them using a bit-copy tool.

Figure 13-2

Using bit copy-tools to capture drive contents



Forensics Toolkit	Comments
Forensic Toolkit (FTK)	Includes a disk imaging program, scanning tools to examine disks, and multiple add-ons for additional capabilities. Developed by AccessData.
Computer Online Forensic Evidence Extractor (COFEE)	Captures and analyzes data from Microsoft systems. Developed by Microsoft and distributed internationally by the International Criminal Police Organization (INTERPOL).
The Coroner's Toolkit	A suite of free tools used to capture and analyze data from UNIX-based systems. Developed by Dan Farmer and Wietse Venema.
Sleuth Kit	A Coroner's Toolkit replacement that captures and analyzes data from UNIX and Windows systems. Developed by Brian Carrier.
EnCase	A widely used suite of products sold by Guidance Software, which also provides EnCase training and certification.
SANS Investigative Forensic Toolkit (SIFT)	A VMware appliance that includes multiple other tools (such as the Sleuth Kit). Developed by the SANS Institute.

Table 13-1 Common Forensics Toolkits

Forensic experts typically use one or more of the available forensic toolkits when collecting, authenticating, and analyzing evidence. These toolkits allow investigators to capture data more efficiently, with less loss and less effort. Table 13-1 shows some of the tools that are available.



TIP Attackers developed the Detect and Eliminate Computer Assisted Forensics (DECAF) tool as a response to COFEE. It attempts to obstruct COFEE's forensic collection methods. However, DECAF does not block any other tools.

Comparing Volatile Data and Persistent Data

Volatile data is data that will be lost if a system is powered down. Data stored in random access memory (RAM) is an example of volatile data. RAM often holds valuable data related to recent activity on a system that can be used in an investigation. However, if the system is powered down before this data is collected, the data is lost. Several computer forensic tools are available that can capture volatile data before the system is shut down. However, there is very little possibility of retrieving information from volatile RAM after power has been removed.

In contrast, some data on a system is *persistent data*, which is data stored on a hard drive or other type of media such as an optical disc or USB flash drive. Persistent data remains even when a system is powered down.



EXAM TIP Systems that need to be investigated should not be powered down until an investigator has had an opportunity to gather data. Forensics tools can capture data in volatile RAM before the system is powered down, but the data is lost once the system is powered down.

Authenticate the Evidence

As soon as evidence is collected, investigators create a chain of custody. A *chain of custody* is a form that documents who obtained the evidence, where it was obtained, when it was obtained, and how it was controlled since it was first collected. Maintaining control of evidence ensures that the evidence is never out of someone's control and subject to tampering. The chain-of-custody form provides proof that the evidence has been controlled since it was collected, and records every interaction with the evidence since it was collected. In the absence of a chain-of-custody form, courts may question the validity of the evidence and refuse to allow its use in legal proceedings.



EXAM TIP A chain-of-custody document provides proof that evidence has been controlled and protected. Without the chain-of-custody document, a court will likely question the validity of the evidence and disallow its use.

In contrast, if an investigator confiscates a laptop as evidence but leaves it on his desk, anyone with access to the desk may be able to add, modify, or corrupt the data on the laptop. Even if no one actually tampered with the laptop while it was on the desk, the laptop may not be admissible as evidence in a court of law without a chain of custody. Even if the investigator locked the laptop in a safe, if he neglected to document his actions with a chain of custody, legal authorities are less likely to believe it was controlled at all times. Controlling the evidence and documenting the control with a chain-of-custody form provides proof that the evidence has been controlled and was not subject to tampering.

You can also authenticate files and drives by comparing a hash value of the original to a hash value of the copy. Chapter 14 covers hashing in more depth, but as an introduction, a hashing algorithm creates a hash, which is simply a number. A hashing algorithm will always create the same hash when executed against the same data, including a forensically sound copy of the same data.

Imagine the hash of the original disk is 123. It should always be 123 no matter how many times you execute the hashing algorithm against the disk. If it is ever anything other than 123, it indicates the original disk has been modified. Similarly, the hash of a forensically sound copy of the disk should be 123. If it's different, it indicates that the copy is not a true bit-by-bit copy of the original disk, and courts may not accept the results from analyzing the copy as admissible evidence.

In reality, a hash is much longer than just three digits. For example, Message Digest 5 (MD5) creates a 128-bit hash expressed as a 32-digit hexadecimal number. However, for the explanation, it's much easier to conceptualize numbers such as 123 rather than numbers like this: dc647eb65e6711e155375218212b3964.

Analyze the Evidence

In the last phase of a computer investigation, the forensic experts analyze the evidence. One of the key requirements in this phase is to use repeatable processes by analyzing copies of the data whenever possible. This ensures the original data is not modified and can be copied again if necessary.

In some cases, it is necessary to analyze the original. For example, magnetic media hard drives often hold data remnants as data remanence. This data remanence is not copied in a bit-by-bit copy but might hold valuable information. The only way to analyze data remanence is by analyzing the original hard drive. This type of analysis should only be done by trained personnel, and only when necessary.

The science of computer forensics has advanced significantly in the last few years. There are many more college and university courses and degrees focused on this science, and there are also several specialized certifications available. Additionally, forensic tools continue to improve, giving forensic experts more and more capabilities.

Forensic Evidence Guidelines and Principles

When an organization decides to add forensic capabilities, it's important to document some basic principles and guidelines within a forensics security policy. These principles and guidelines can be anything the organization chooses to add, and many professional organizations have provided suggestions of items to include.

The Association of Chief Police Officers (ACPO), a British organization recently replaced by the National Police Chiefs' Council (NPCC), published the *ACPO Good Practice Guide for Digital Evidence*. Version 5.0, published in 2012, includes the following four principles:

- **Principle 1** No action taken by law enforcement agencies, persons employed within those agencies or their agents should change data which may subsequently be relied upon in court.
- **Principle 2** In circumstances where a person finds it necessary to access original data, that person must be competent to do so and be able to give evidence explaining the relevance and the implications of their actions.
- **Principle 3** An audit trail or other record of all processes applied to digital evidence should be created and preserved. An independent third party should be able to examine those processes and achieve the same result.
- **Principle 4** The person in charge of the investigation has overall responsibility for ensuring that the law and these principles are adhered to.

NIST SP 800-86, *Guide to Integrating Forensic Techniques into Incident Response*, stresses the importance of clearly defining the roles and responsibilities of personnel performing forensic activities. It mentions that forensic processes should be consistent with the organization's policies and all applicable laws. It also stresses the importance of maintaining a chain of custody to support the admissibility of evidence into legal proceedings.

NIST SP 800-72, *Guidelines on PDA Forensics*, suggests using the Daubert method, which is a set of standards that help ensure the reliability of evidence. When applied to digital evidence, it's important to consider the following reliability factors derived from the Daubert method and set forth in SP 800-72. Following these reliability factors will help ensure the evidence can be used in a court of law.

- **Testability** Has the scientific theory or technique been empirically tested? (Some tools, such as the EnCase Toolkit, have been used successfully in court cases, but some have not.)
- **Acceptance** Has the scientific theory or technique been subjected to peer review and publication? This ensures that flaws in the methodology would have been detected and that the technique is finding its way into use via the literature.
- **Error Rate** What is the known or potential error rate? Scientific measures generally have associated error rates, which can be estimated with a fair amount of precision.
- **Credibility** What is the expert's qualifications and stature in the scientific community? Does the technique rely upon the special skills and equipment of one expert, or can it be replicated by other experts elsewhere?
- **Clarity** Can the technique and its results be explained with sufficient clarity and simplicity so that the court and the jury can understand its plain meaning?



TIP It's rarely clear at the beginning of an investigation that evidence will be needed in a court of law. Because of this, personnel often overlook important evidence, handle it improperly, or accidentally destroy it early in an investigation. However, by always following proper forensic practices, it ensures the protection of evidence in case it is ever needed.

Comparing Computer Abuse and Computer Crime

Computer abuse and computer crime mean two separate things. In computer abuse, an individual is dishonest or unethical and using a computer in an unauthorized manner. In contrast, a criminal commits illegal acts with a computer. Both of these concepts are highlighted in Figure 13-3.

Figure 13-3
Computer abuse and computer crime

Computer Abuse	Computer Crime
Counter to established policies	Breaks a law
Misuse Unethical use	Theft Destruction

For example, in computer abuse, an individual may misuse a computer system in violation of an organization's security policies. The user may be using the computer to shop for personal items during company time despite knowing that this is against the established acceptable use policy.

Cybercrime (or computer crime) is any crime that involves a computer. The criminal may use a computer as a tool to commit the crime, or the computer may be the target of the crime. The purpose of cybercrime is often monetary gain for the criminal, but some criminal acts are simply destructive, with the purpose of negatively affecting a computer or organization.



EXAM TIP Computer abuse may result in the loss of a job or position, but does not actually involve breaking laws. Computer crime includes illegal activities and can result in prosecution, fines, and/or jail time.

The three commonly quoted elements needed to commit a crime are means, motive, and opportunity (sometimes referred to as means, opportunity, and motive, or MOM). By understanding these, you have a greater chance at reducing crime. More specifically, if you can remove the means, motive, or opportunity to commit a crime, you have a better chance of preventing the crime. The means is the ability of the criminal to commit the crime, the motive is the reason (such as for money or revenge), and the opportunity is simply the chance to commit the crime.

The following are some of the common terms related to computer abuse and computer crimes:

- **Cyberbullying** Cyberbullying has become common among some young people. The bully uses the Internet to attack a victim through harassment, coercion, or intimidation. For example, a cyberbully may spread lies about a person via social media such as Facebook or Twitter, or send harassing text messages to the victim or to others about the victim.
- **Cyberstalking** Cyberstalking includes cyberbullying actions, but it goes further. For example, the stalker may gather information on the victim, monitor the victim's activities, and steal the victim's identity. Adults are typically involved with cyberstalking, while children are more likely to be involved in cyberbullying. A number of laws are in place addressing cyberstalking making it a criminal act.
- **Cyberwarfare** Cyberwarfare is a politically motivated attack on entities in another country. The goal is typically sabotage to disrupt the capabilities of an entity in the target country or espionage to gather intelligence on the target. For example, the U.S. government concluded that North Korea attacked Sony in late 2014 and labeled this cyberwarfare. Similarly, the *Washington Post* reported that attackers accessed computer networks in the White House and that the *Post's* sources indicated the attackers were Russian state-backed groups. Attack groups sponsored by a state (or government) are typically referred to as advanced persistent threats (APTs).

Understanding Fraud and Embezzlement Crime

Fraud is the use of deception for unlawful gain or unjust advantage. It includes engaging in wrongful or illegal activities with the intention of gaining undue or illegitimate advantage over another person, group, or entity. Embezzlement is a specific type of financial fraud where an individual steals money or property from an employer, company, clients, customers, or even partners. It also includes misappropriating money or assets held in trust.

MCI WorldCom gained notoriety for a massive accounting scandal that resulted from fraudulent accounting methods. CEO Bernard Ebbers helped grow the business that ultimately became MCI WorldCom. Later, courts held him responsible for the company's fraudulent financial reporting that ultimately resulted in the loss of over \$100 billion to investors, and he was ultimately convicted of fraud and conspiracy and is serving a 25-year prison sentence.



TIP Chapter 11 mentions the Sarbanes-Oxley (SOX) Act of 2002. It requires CEOs and other executives to personally verify the accuracy of financial data. Part of the goal of the legislation is to prevent fraudulent financial reporting similar to what occurred with MCI WorldCom.

Bernie Madoff has become a household name related to fraud. He managed a Ponzi scheme for years under the guise of an investment company. Many people regarded him as an outstanding investment banker who ran a successful wealth management company. In reality, he routinely skimmed new investments and used these investments to pay what looked like great dividends to current investors. Investigators calculated the amount of missing funds at about \$65 billion. He is currently serving a 150-year prison term. Interestingly, investigations such as one by the U.S. Securities and Exchange Commission failed to uncover the fraud, which went undetected until Madoff admitted the Ponzi scheme to his sons in 2008.

Chapter 5 mentions salami attacks where an attacker takes multiple minor actions to get large gains. For example, prosecutors indicted Michael Largent in 2008 for collecting more than \$50,000 from some online brokerage firms by collecting between a penny and \$2 at a time. He used scripts to open over 58,000 accounts under phony names with the brokerage firms. These firms verified his bank account by depositing small amounts of money into his account. However, when his bank questioned the thousands of small payments, he didn't have any good answers and the FBI later arrested him.

Many companies use job rotation and mandatory vacations to reduce the risks of internal fraud and embezzlement. These policies also reduce illegal data modifications and general misuse of information.

Mandatory Vacations

Many banking institutions implement mandatory vacation policies. These require employees to take a vacation outside of the workplace for a specified period, such as at least five or ten consecutive workdays. The goal is to require someone else to perform their job functions, which increases the possibility of exposing any suspicious activities.

For example, imagine that an employee identifies a loophole in the flow of money within an organization and discovers that he can divert a small percentage of money in each transaction. However, if he ever stopped diverting the money, the loss would appear as an anomaly. If the organization forces him to take time off work, the diversion would stop for a short time, and someone would discover the fraud. Knowing this is a possibility will often deter an employee from exploiting the loophole in the first place.

Similarly, imagine if an investment banker takes millions of dollars from specific investors but never actually invests the money. Instead, he uses the money to pay healthy dividends to other investors who previously gave him money. He gives these investors fake documents indicating that they have a substantial amount of wealth. In reality, this crooked investment banker has stolen all the money. As long as the investment banker is at work, he can answer the queries from the investors and no one will discover his actions. However, if the organization forces him to take regular vacations, the Ponzi scheme will collapse while he is away.

Admittedly, a mandatory vacation policy doesn't guarantee that fraud won't occur. Consider this same scenario with two or more people colluding to create and manage this Ponzi scheme. If that's the case, forced vacations won't necessarily discover the actions. The accomplices can cover for anyone who is on vacation. However, when two or more people are involved, it increases the chances that someone will make a mistake. Additionally, law-breaking criminals committing fraud aren't necessarily trustworthy, so any one of them might report the actions at some point and place all the blame on one of the other accomplices.

While mandatory vacations can be an effective security control, they aren't foolproof. However, when used as part of an overall defense-in-depth plan, mandatory vacations can help reduce the chance of fraud and embezzlement.



EXAM TIP A mandatory vacation policy requires employees to take vacations for a minimum period (such as five or ten consecutive days). The goal of the policy is to reduce the chance of fraud.

Job Rotation

Job rotation is the practice of moving employees periodically, such as once a year or so, between different jobs. The goal is for employees to do different jobs at different times, which also prevents any single employee from doing the same job for years at a time.

Chapter 1 covered separation of duties and explained how it prevents a single person from having complete control over a process. For example, accounting departments frequently separate accounts payable tasks from accounts receivable tasks to reduce the risks. Even though a separation of duties policy separates these tasks, employees working in an accounting department typically have the skills to do both jobs. In other words, someone working in an accounts payable division also has the skills to perform accounts receivable tasks. A job rotation policy periodically moves personnel between the two jobs. Just in case an employee discovers a loophole in the process that he or she can use to defraud the company, the job rotation policy limits how long the fraud could continue. It also increases the chance of discovering the fraud, which acts as a deterrent.

Many organizations use job rotation for cross-training. For example, organizations often rotate management personnel through different jobs within an organization to gain

a better overall understanding of the organization. The primary purpose of this practice is to help build the knowledge and skills of personnel rather than deter and detect fraud. Job rotation also helps in disaster recovery if key personnel are unavailable.



EXAM TIP Job rotation used as a security practice can reduce the risk of fraud within a company. It is commonly combined with the security practice of separation of duties.

Understanding Privacy Issues

Although the explosion of the use of computers has certainly helped people share information, it has also resulted in people having easy access to massive amounts of personal data about individuals. If organizations don't protect this data, it can negatively affect an individual's privacy. Even if individuals are very careful with their privacy data and reveal it only when required to do so for business transactions, they are still affected if the business suffers a data breach. If these businesses don't protect the customer's data, the customer loses.

Because of this, many laws require organizations to protect personal data. Despite these laws, data breaches continue. For example, the Identity Theft Resource Center (ITRC) tracks reported data breaches and reported an average of 15 data breaches a week in 2014. ITRC's 2014 report identified 783 reported data breaches exposing data from more than 85 million records. Some of the categories of exposed data include

- **Business** 258 data breaches exposing more than 68 million records
- **Government/military** 92 data breaches exposing 6.6 million records
- **Medical/health care** 333 data breaches exposing more than 8 million records
- **Banking/credit/financial** 43 data breaches exposing almost 2 million records

Legal experts agree that the key to avoiding liability for any data leakage is due diligence. In other words, if an organization is actively taking steps to protect data and prevent privacy compromises, but still suffers a loss, it is less likely to be held liable. The message seems to be that "incidents happen," but as long as you're doing your best to prevent the incidents, you'll have less legal liability.

Chapter 11 introduced personally identifiable information (PII) and the importance of protecting it. PII is information that can be used to personally identify an individual and includes components such as those shown in Figure 13-4. Chapter 11 also introduced the Health Insurance Portability and Accountability Act (HIPAA) of 1996. HIPAA mandates the protection of protected health information (PHI), which is any PII that includes information concerning the health status, provision of health care, or payment of health care for an individual.



TIP You don't need to have all of the components shown in Figure 13-4 to have PII. If any single piece of information combined with a name can uniquely identify an individual, it is PII.

The following sections cover some of the regulations that mandate the protection of PII.

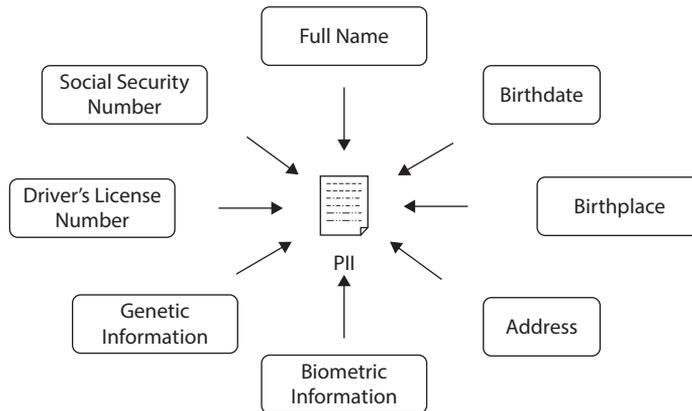


Figure 13-4 Components of PII

Defining a Data Breach

Generically, a data breach is any incident where unauthorized personnel gain access to sensitive, protected, or confidential data. It often occurs when criminals from outside an organization launch attacks allowing them to access data held inside an organization.

The Identity Theft Resource Center focuses on the exposure of personal data and defines a data breach as “an incident in which an individual name plus a Social Security number, driver’s license number, medical record or financial record (credit/debit cards included) is potentially put at risk because of exposure. This exposure can occur either electronically or in paper format.” Outside of the United States, this definition of a data breach includes a national identification number instead of a Social Security number. Note that this matches the definition of PII. Because medical records are included, it also matches the definition for PHI.

Some data breaches only expose usernames, passwords, and e-mail addresses without exposing other information that identifies an individual. The ITRC doesn’t report the number of records exposing only this data, primarily because its focus is on identity theft. However, many other definitions include this data compromise as a data breach. Attackers can do a lot of damage if they gain access to valid usernames, passwords, and e-mail addresses.

Similarly, the security firm Mandiant discussed the theft of intellectual property in its 2013 report titled *APT1: Exposing One of China’s Cyber Espionage Units*. In one instance, Mandiant observed APT1 stealing 6.5 terabytes of data from a single organization over a ten-month time period. This might not have included personal data, so it doesn’t meet the ITRC definition of a data breach. However, security professionals, working to protect all data, consider this a serious data breach.

European Directives

The European Union (EU) has enacted several directives to protect privacy. Two of these are the Data Protection Directive (Directive 95/46/EC) and the E-Privacy Directive (Directive 2002/58).

The Data Protection Directive regulates the processing and protection of personal data within the EU. This directive defines personal data as “any information relating to an identified or identifiable natural person (‘data subject’),” which is very similar to the definition of PII. It mandates specific steps to protect personal data. Further, it restricts data transfers to countries outside of the EU.



TIP The European General Data Protection Regulation is in draft stage as I write this, but it is expected to replace the Data Protection Directive. EU officials expect to have it finalized and adopted in 2015 or 2016, and it will take effect two years after adoption.

The U.S. Department of Commerce worked with EU officials and developed the Safe Harbor program. It helps U.S. organizations comply with the requirements of the Data Protection Directive. The Safe Harbor program includes the following seven principles. U.S. organizations that agree to follow these principles can transfer data to and from EU organizations and will remain in compliance with the Data Protection Directive as long as they adhere to these principles.

- **Notice** Organizations must inform individuals about what data they are collecting and how the organization plans to use the data.
- **Choice** Individuals must be able to opt out of the data collection.
- **Onward Transfer** Organizations agree to transfer the data only to other organizations that comply with the notice and choice principles.
- **Security** Organizations must take reasonable precautions to protect the data.
- **Data Integrity** Organizations will only collect data that is relevant to the stated purpose in the notice and agreed to by the individual.
- **Access** Individuals must be able to access their personal data and have the ability to correct, amend, or delete information if it is inaccurate.
- **Enforcement** Organizations must implement mechanisms to assure compliance with the principles.

The E-Privacy Directive focuses on the protection of digital data and includes regulations concerning the transmission of data. Directive 2009/136, which defines and regulates the use of cookies, amended this directive. Organizations often use cookies to collect data on website visitors and use this data for analytics. The E-Privacy Directive requires organizations to inform users of how the cookies are used and give users the ability to opt out of data collection methods.



TIP Many sites state that you can disable cookies. However, doing so also disables a lot of site usability. Developers commonly use cookies to maintain the state of user interaction with a site. When cookies are disabled, the site cannot maintain the user state from one page to another, which can significantly degrade the user experience.

California Supreme Court Rules That ZIP Codes Are PII

The California Supreme Court determined in early 2011 that a person's Zone Improvement Plan (ZIP) code is PII. In this case, Williams-Sonoma stores collected a customer's ZIP code each time that the customer made a credit card purchase. It then combined the customer's name and ZIP code in a search to discover the customer's address (which the customer didn't provide) and used this information for marketing purposes.

The court determined that a ZIP code is not required for a credit card transaction, but is readily understood to be part of an address. Because a company can combine a ZIP code with the person's name to discover the person's address, the California Supreme Court classified it as PII. The decision restricts businesses in California from requesting a ZIP code during a credit card transaction. The California Supreme Court's justification in this case is that the business can use this personal information for business purposes.

While the California Supreme Court determined that a ZIP code is PII when combined with a customer name, it is one of the few jurisdictions to do so. If you buy gas with a credit card in a different jurisdiction, you are often required to enter the ZIP code.



EXAM TIP Data inference (mentioned in Chapter 11) is the ability to piece together knowledge with seemingly unconnected information. In this case, the company was using a data inference technique to get a customer's address from online databases.

Connecticut's Public Act No. 08-167

In 2008, Connecticut passed a law titled "An Act Concerning the Confidentiality of Social Security Numbers." Although the title focuses on Social Security numbers, the contents of this law go beyond just these numbers.

For example, it states, "Any person in possession of personal information of another person shall safeguard the data, computer files, and documents containing the information from misuse by third parties..." It then defines personal information as "information capable of being associated with a particular individual through one or more identifiers, including, but not limited to, a Social Security number, a driver's license number, a state identification card number, an account number, a credit or debit card number, a passport number, an alien registration number or a health insurance identification number..."

Intentional violations of this law can result in a fine of \$500 for each violation, with a maximum fine of \$500,000 per event. In other words, if a database is released that includes this information on 1,000 or more people, it can result in a fine of half a million dollars.

A notable distinction with this law is that unintentional violations are not subject to the law. In this case, if an organization discloses the information, it must prove that the disclosure wasn't intentional. The law does not consider a data breach by an attacker an intentional violation as long as the organization took steps to protect the information.

Children's Online Privacy Protection Act

The Children's Online Privacy Protection Act (COPPA) is a U.S. federal law that applies to the collection of information on children under the age of 13. The goal is to protect the privacy of these children by preventing websites or online services from collecting information on them without parental consent.

For example, UMG Recordings was fined \$400,000 after it was found to have knowingly collected information on children under the age of 13. The company created a website directed at children to promote a 13-year-old pop star known as Lil' Romeo. UMG Recordings collected extensive information on the children, including their full name, e-mail address, home address, phone number, interests, and birthdate. Even when the user's birthdate clearly indicated that the user was under the age of 13, UMG Recordings still collected and used the information. For example, it used the collected e-mail addresses to send e-mail marketing about other musicians and other websites.

Although it is possible to collect information on children under 13, the organization must gain parental consent, which is often challenging. Because of this, most organizations simply do not collect information on these children to remain in compliance of the law.

California Online Privacy Protection Act of 2003

The California Online Privacy Protection Act (OPPA) of 2003 became law in 2004. Notice the acronym is OPPA without an initial *C* to differentiate it from the Children's Online Privacy Protection Act (COPPA). It requires operators of commercial websites to conspicuously post and comply with a privacy policy if they collect PII on any visitors who reside in California. In this context, a commercial website is one created by or for an organization. In other words, OPPA doesn't apply to personal websites created by or for individuals.

Note that this law doesn't just apply to websites created by organizations in California or hosted on servers in California. Instead, it applies to any organization that makes its website available to any California residents. In other words, it applies to any website on the Internet, because California residents with Internet access can access pages hosted anywhere in the world.

Although it's doubtful that California law will extend to remote places like Tanzania, it will certainly affect websites operating within the United States. Because of this, most operators of commercial websites in the United States make certain that they include a privacy policy that complies with this law if they collect any information on the user, such as name, e-mail address, or other data, used to track the user.

The privacy policy must be easy to locate (for example, websites typically make it accessible from the home page) and include the following:

- What information is gathered by the website
- How the information is used and shared
- The effective date of the policy
- A description of changes since the policy was first created

Chapter Review

Computer forensics is the science of examining and inspecting computer systems for evidence about an event or crime. IT and security personnel will often participate in incident response, so they should understand the phases of incident handling. This includes discovering an incident, responding to an incident, escalating an incident, reporting an incident, implementing feedback loops (or lessons learned reviews), and implementing countermeasures.

First responders are the first forensics-trained personnel on the scene, and one of their primary responsibilities is the preservation of the scene. The three phases of a typical computer forensic investigation are acquiring the evidence, authenticating the evidence, and analyzing the evidence.

When acquiring the evidence, one of the most important concerns is to protect the data and prevent data modifications. Forensic tools allow experts to capture data in volatile RAM before a system is powered down and use bit-copying procedures to copy the contents of hard drives. Several forensic tools are available such as FTK, COFEE, the Sleuth Kit, EnCase, and more.

In the authentication stage, a chain-of-custody document is used to provide proof that the evidence was controlled after it was collected. The chain-of-custody document ensures that the evidence is admissible in a court of law. The analysis of the evidence is the last stage. Some of the same forensic tool suites used to collect evidence include tools that can be used to analyze evidence.

Computer abuse indicates that an individual is using a computer or network in an unauthorized manner, such as counter to an organization's security policies. Computer crime indicates that the user is committing illegal acts with the computer. Computer crime is punishable by fines and/or jail time.

Computers can be used in fraud and embezzlement crimes resulting in substantial losses to an organization. Mandatory vacations and job rotation are two security policies that an organization can use to reduce the success of fraud and embezzlement crimes. These two policies are often used with other security controls as part of an overall defense-in-depth strategy.

PII includes any information that can identify an individual. Privacy issues related to the protection of PII have become more important as the amount of data that organizations collect on their customers has increased. If an organization doesn't take precautions to protect customer data, attackers may breach the organization's database and access its customer data. Many laws have been enacted on the international, national, and state levels requiring organizations to protect PII.

Questions

1. Which of the following *best* describes a primary goal of incident handling?
 - A. Collecting evidence
 - B. Documenting evidence
 - C. Containing any potential damage
 - D. Improving security controls
2. Which one of the following is mostly like to be performed during a feedback loop in an incident handling process?
 - A. Chain of custody
 - B. Hashing
 - C. Escalation
 - D. Perform a lessons learned review
3. Which of the following identifies a primary responsibility of a first responder after a computer security incident?
 - A. Capturing images of disks
 - B. Capturing data in RAM
 - C. Interviewing witnesses
 - D. Preserving the scene
4. What forensic evidence can be lost if a system is powered down before the evidence is collected?
 - A. Data on the disk drive
 - B. Data on a USB drive
 - C. Data in memory
 - D. Data in files
5. Managers suspect that an employee in your organization has committed fraud. You are told to secure his computer as part of an incident response. Which of the following should you NOT do?
 - A. Disconnect the computer from the LAN
 - B. Power the system down
 - C. Prevent anyone from accessing the system
 - D. Take pictures of the system
6. Which of the following is a benefit of a chain-of-custody form?
 - A. It helps ensure that evidence is protected.
 - B. It helps ensure that evidence is controlled.
 - C. It helps ensure that evidence is not modified.
 - D. It helps ensure that evidence is admissible in court.

7. Of the following choices, what is NOT a phase of a computer forensic investigation?
 - A. Prosecution based on evidence
 - B. Authenticating evidence
 - C. Analyzing evidence
 - D. Acquiring evidence
8. Which of the following choices is the most important consideration when gathering evidence as part of a computer forensic investigation?
 - A. Ensuring that systems are turned off as soon as possible
 - B. Ensuring that a record of files on a system is recorded by accessing the system
 - C. Ensuring that users can log on to the system
 - D. Ensuring that evidence is not modified
9. A forensic expert wants to examine data on a hard drive of a confiscated computer. Which of the following actions should the expert complete first?
 - A. Ensure that the computer has UPS protection
 - B. Create a bit copy of the disk
 - C. Disable the antivirus software on the computer
 - D. Move the hard drive to another system and examine it on the other system
10. Which of the following helps to prove that collected evidence has been controlled since it was collected?
 - A. COFEE application
 - B. Chain of custody
 - C. DECAF application
 - D. Audit logs
11. In general, what elements need to come together for a crime?
 - A. Means, motive, and opportunity
 - B. Criminal, software, and hardware
 - C. Discovery, theft, and benefit
 - D. Attacker, attackee, and method
12. Countries sometimes engage in espionage against other countries. What is this called?
 - A. Cyberbullying
 - B. Cyberstalking
 - C. Cyberwarfare
 - D. Cyberterrorism

13. What is the purpose of mandatory vacations in relation to security?
 - A. To ensure that employees do not burn out
 - B. To ensure that employees take time to relax
 - C. To reduce the payroll of an organization
 - D. To reduce the chance of fraud
14. What security practice moves employees into different positions periodically to reduce the risk of fraud?
 - A. Separation of duties
 - B. Job rotation
 - C. Mandatory vacations
 - D. Risk mitigation
15. A business in Florida gathers customers' names and ZIP codes and uses them to identify the customers' addresses. What is occurring?
 - A. Violation of an EU directive
 - B. Data breach
 - C. Data inference
 - D. Violation of COPPA
16. An organization collects customer data such as their name, e-mail address, physical address, and phone number. What term best describes this information?
 - A. PHI
 - B. PII
 - C. COFEE
 - D. DECAF
17. An organization regularly collects information on customers for marketing purposes. It uses this information to personally identify the customers. Who is responsible for protecting this data?
 - A. The organization.
 - B. It depends on whether the customers gave permission to collect the data.
 - C. It depends on whether a data breach occurred.
 - D. Customers.
18. What are the principles of notice, choice, access, and enforcement most closely related to?
 - A. Privacy policies
 - B. Incident response
 - C. Safe Harbor
 - D. Protection of children's privacy

19. What law requires an organization to get a parent's consent prior to collecting information on children under 13?
 - A. COPPA
 - B. OPPIA
 - C. Data Protection Directive
 - D. E-Privacy Directive
20. What law requires organizations to post a privacy policy on their website?
 - A. SOX
 - B. PHI
 - C. OPPIA
 - D. COPPA

Answers

1. **C.** Primary goals of incident handling include containing any potential damage and repairing the damage that has occurred. While incident handling often includes evidence collection and using chain-of-custody forms to document collected evidence, these are not primary goals. The feedback loop (or lessons learned) process often improves security controls, but this is not a primary goal.
2. **D.** Feedback loops in an incident handling process include a lessons learned review, which examines the incident and the response. The goal is to prevent a recurrence of the incident. A chain of custody validates that evidence has been controlled since it was collected. Hashing provides assurances that data (such as a bit-by-bit copy of a disk) is identical. Escalation refers to getting more people involved in an incident.
3. **D.** A primary responsibility of first responders is the preservation of the scene. Forensics experts acquire and analyze the evidence, which includes capturing images of disks and data in RAM. Investigators interview witnesses.
4. **C.** Data in memory (volatile RAM) is lost if a system is powered down. Power isn't required to retain data on disk drives or USB drives, so it is not lost when powered down. Data in files is retained on disk drives.
5. **B.** You should not power the system down, as doing so will delete any data in volatile RAM. All of the other answers are valid actions in response to an incident that may result in legal proceedings. Organizations often have policies to disconnect a computer from the local area network (LAN) to isolate it and protect evidence on the system. No one should be allowed to access the system as a precaution to ensure that data (and potential evidence) is not modified. There's nothing wrong with taking pictures of the system to preserve evidence.

6. **D.** A chain-of-custody form provides proof that evidence has been protected and helps ensure that the evidence is admissible in court. Note that it only documents how the evidence has been protected and controlled, but it doesn't actually protect or control the evidence. It also doesn't ensure that the evidence is not modified.
7. **A.** Prosecution is not a part of a computer forensic investigation, but may occur based on the results of an investigation. A computer forensic investigation typically includes the three phases of acquiring, authenticating, and analyzing evidence.
8. **D.** It's important that evidence is not modified when collecting evidence. If a system is turned off, data in volatile RAM is lost. If files on a system are accessed, it modifies the files by showing new access times. If a user logs on to a system, it can modify information about the last user that logged on.
9. **B.** Disk drives should be copied with an approved bit-copy tool prior to examining any of the data because the process of examining the data can result in the loss of data. An uninterruptible power supply (UPS) protects a system from a power loss, and antivirus (AV) software protects a system from malware, but neither has anything to do with protecting evidence. Moving the drive to another system and examining it on the other system will modify the original evidence.
10. **B.** A chain-of-custody document provides proof that collected evidence has been controlled since it was collected by identifying exactly who has control of it and where it was located since it was collected. Microsoft developed the Computer Online Forensic Evidence Extractor (COFEE) to help law enforcement forensic experts, and hackers developed Detect and Eliminate Computer Assisted Forensics (DECAF) as an antiforensics tool. Neither COFEE nor DECAF provides proof that the evidence was controlled after collection. Audit logs can be used in a forensic investigation but they don't validate evidence.
11. **A.** The three commonly quoted elements for a crime are means (the ability to commit the crime), motive (such as money or revenge), and opportunity (the chance to commit the crime). Crimes can be committed without software or hardware. Criminals can commit crimes (such as vandalism or destruction) without theft. Similarly, criminals can commit crimes (such as theft) without an attack.
12. **C.** Cyberwarfare is a politically motivated attack on entities in another country and is done for sabotage and/or espionage. Cyberbullying occurs when one person harasses, coerces, or intimidates another person using the Internet. Cyberstalking is more serious than cyberbullying and is a criminal act. Cyberterrorism is the use of the Internet to launch terrorist attacks.
13. **D.** A mandatory vacation policy can reduce the chance of fraud by requiring other employees to take over the tasks and responsibilities of a vacationing employee. While vacations are good to help employees relax and reduce the chance of burnout, these matters aren't as relevant to security issues as reducing the chance of fraud. While some companies do implement mandatory vacations to reduce payroll expenses during periods of low activity, this is not a security policy.

14. **B.** Job rotation is a security practice of rotating employees into different positions periodically to reduce the risk of fraud. It is often combined with a separation of duties policy, but separation of duties by itself does not require employees to rotate between jobs; only that a single employee doesn't control all elements of a process. A mandatory vacation security policy requires employees to take time away from work to increase the chances of suspicious activity being discovered. Risk mitigation is any practice that reduces risks, but it is not limited to only rotating employees to different positions or only reducing the risk of fraud.
15. **C.** This is an example of data inference, because the company is collecting small pieces of information to get other information that isn't provided directly. Since it's occurring in Florida, it isn't in violation of a European Union (EU) directive. A data breach occurs when unauthorized individuals access stored data, but there's no indication of unauthorized access in this scenario. The Children's Online Privacy Protection Act (COPPA) protects the privacy of children under 13.
16. **B.** This information is personally identifiable information (PII) because it can be used to identify the customers personally. Protected health information (PHI) is information concerning the health status, provision of health care, or payment of health care for an individual. COFEE is a forensic tool used by law enforcement agencies, and DECAF is an antiforensics tool designed to detect COFEE.
17. **A.** Organizations are responsible for protecting personally identifiable information (PII), which is information that can identify individuals. This is true even if the customers give permission to the organization to collect the data, and even if a data breach has not occurred. Customers are not required to protect the information that they give to an organization.
18. **C.** These are four of the seven principles in the Safe Harbor program. This program helps organizations protect data transferred to and from the European Union (EU). The California Online Privacy Protection Act of 2003 (OPPA) mandates the use of privacy policies for websites viewable in California, but it doesn't mention these principles. These principles are not related to incident response. The Children's Online Privacy Protection Act (COPPA) focuses on the protection of children's privacy, but it doesn't mention these principles.
19. **A.** The Children's Online Privacy Protection Act (COPPA) requires organizations to get a parent's consent prior to collecting information on children under 13. The California Online Privacy Protection Act of 2003 (OPPA) requires operators of commercial websites to post a privacy policy on the website if the website collects personally identifiable information (PII). The Data Protection Directive (Directive 95/46/EC) restricts data transfers of privacy data to countries outside of the European Union. The E-Privacy Directive (European Directive 95/46/EC) focuses on the protection of digital data and regulates the use of cookies.
20. **C.** The California Online Privacy Protection Act of 2003 (OPPA) requires operators of commercial websites to post a privacy policy on the website and comply with it if the website collects personally identifiable information (PII).

Sarbanes-Oxley (SOX) Act is a U.S. law that requires executives of publicly held companies to validate the integrity of their financial data. HIPAA is a law focused on the protection of protected health information (PHI) and requires organizations to post a privacy policy. However, PHI is not a law. The Children's Online Privacy Protection Act (COPPA) is a U.S. federal law that applies to the collection of information on children under the age of 13.

Cryptography

In this chapter, you will learn about

- Basic cryptography concepts and terms
- Hashing algorithms and integrity concepts
- Common symmetric encryption algorithms
- Differences between symmetric and asymmetric encryption
- Asymmetric encryption using public and private keys
- Asymmetric encryption and digital signatures
- Hiding data within data using steganography
- The components of a public key infrastructure
- Different types of cryptanalysis attacks

Understanding Basic Cryptography Concepts

Chapter 1 introduces confidentiality and integrity (in the context of the CIA security triad of confidentiality, integrity, and availability). Cryptography includes confidentiality and integrity, but it's worth stressing that many cryptography processes provide authenticity too. Confidentiality from access controls starts with proper identification and authentication as described in Chapter 2. Appropriate access controls then grant personnel access to data based on their proven identity. As an introduction for this chapter, the following list describes confidentiality, integrity, and authenticity within the context of cryptography:

- **Confidentiality** The goal of confidentiality is to prevent unauthorized access to data. Beyond access controls, the primary method of preventing loss of confidentiality is with encryption. This chapter focuses on encryption techniques to protect the confidentiality of data. Two common encryption techniques are symmetric encryption (with one key) and asymmetric encryption (with two keys).
- **Integrity** The goal of integrity is to prevent unauthorized modifications to data and systems. Beyond access controls, hashing is one of the primary methods of providing assurances that data has not lost integrity. A hash is calculated on data at different times, such as on Monday and then again on Tuesday. If the two hashes are the same, it proves that the data is the same and has maintained integrity. This chapter describes hashing and covers several hashing protocols.

- **Authenticity** Validation processes provide assurances that data is coming from a known source and is valid or reliable. Many encryption and hashing protocols provide authenticity. Additionally, some processes, such as digital signatures, have authenticity as a key goal. This chapter covers all of these concepts.



EXAM TIP Cryptography concepts have a lot of depth. However, the SSCP exam doesn't expect you to design cryptographic algorithms or be able to describe the inner workings of them. You do need to have some basic knowledge so that you can recognize differences between some basic concepts.

Cryptography Terminology

Cryptography has many terms that might not be familiar to you. The following list introduces some of these terms and topics, all of which are covered in greater depth in this chapter:

- **Plaintext** Data that is not in an encrypted format. Plaintext data (sometimes called cleartext data) is easily readable by anyone. You can think of it as ordinary text, such as what you're reading on this page.
- **Ciphertext** Data that is in an encrypted, or scrambled, format. Authorized personnel can decrypt ciphertext to return it to plaintext and read it. When strong encryption protocols are used, unauthorized personnel cannot read ciphertext.
- **Encryption** The process of converting plaintext data into ciphertext data to prevent loss of confidentiality. Figure 14-1 shows the overall process of encrypting plaintext data into ciphertext data using an encryption algorithm. Decryption reverses the process by decrypting the ciphertext data with a decryption algorithm to re-create the original plaintext data.

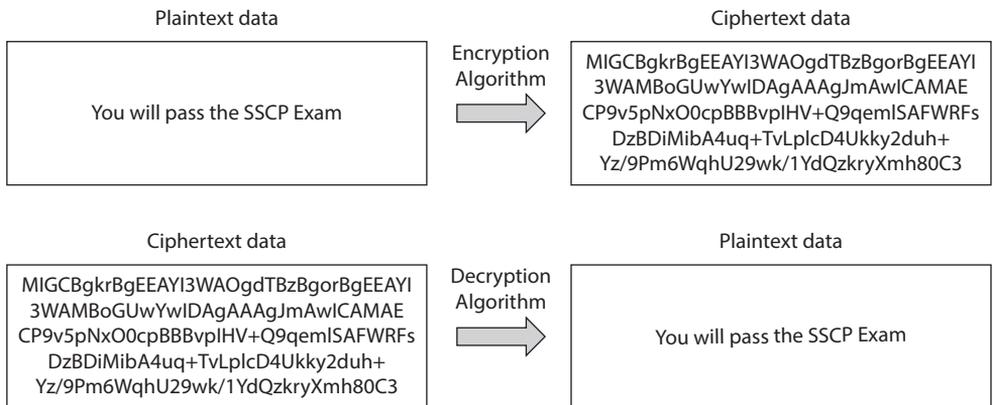


Figure 14-1 Encryption and decryption process

- **Encryption algorithm** A mathematical process that converts plaintext data to ciphertext data. Encryption algorithms include a process to decrypt the data. Encryption algorithms are often published so that anyone can see how they work. For example, you can find the algorithm for Advanced Encryption Standard (AES) by doing an Internet search on “Advanced Encryption Standard algorithm” or “AES algorithm.”
- **Symmetric encryption** Data is encrypted and decrypted with the same key. For example, if an encryption algorithm uses a key of 101 to encrypt data, it will use the same key of 101 to decrypt the data. The algorithm would use a different key (such as 234) to encrypt and decrypt a different set of data. In practice keys are much longer than three digits.
- **Asymmetric encryption** Data is encrypted and decrypted with two keys. These two keys are a matched pair known as a public key and a private key. A private key can decrypt data that has been encrypted with the matching public key. Similarly, a public key can decrypt data that has been encrypted with the matching private key. The private key is always kept private, and the public key is freely shared.



EXAM TIP Encryption helps prevent loss of confidentiality. The two primary categories of encryption are symmetric (using one key) and asymmetric (using two keys, known as a public key and a private key).

- **Cryptography** The science of scrambling, or encrypting, data. Cryptography includes the study of algorithms and other methods that help protect the confidentiality of data.
- **Cryptographic system** A system that includes the algorithms needed to support a particular method of encryption or decryption. Although most systems support standard algorithms, administrators sometimes need to add features to a computer to support a newer algorithm.
- **Cryptanalysis** The science of studying cryptographic methods. Security experts and attackers often use cryptanalysis methods to discover weaknesses in existing algorithms. Attackers attempt to use cryptanalysis to defeat a cryptosystem with the ultimate goal of compromising the confidentiality or integrity of protected data.
- **Hashing** The process of creating a hash from a file or a message with a hashing algorithm. A hash is simply a number. A hashing algorithm will always produce the same hash when executed against the same data. In other words, as long as a file or message never changes, a recalculated hash will always be the same no matter when it is calculated.

Data Sensitivity

The primary purpose of encrypting data is to preserve its confidentiality. However, this isn't required for all data. For example, Chapter 11 discusses data classifications such as

Top Secret, Secret, Confidential, and Unclassified. Data that is classified as Top Secret, Secret, or Confidential is sensitive, so it's appropriate to protect this data with encryption protocols. However, Unclassified data isn't sensitive, so it's often a waste of processing power to encrypt and decrypt unclassified data.

Additionally, the data sensitivity drives the selection of the encryption protocols. Higher data classifications (such as Top Secret) warrant the use of the strongest encryption protocols. In contrast, lower data classifications (such as Confidential) do not need the strongest encryption protocols.

You might be wondering, "Why not use the strongest encryption protocols for everything?" The reason is the performance cost. It takes processing power to encrypt and decrypt data, and if all data is encrypted and decrypted, it can slow the system down. Also, encrypted data typically consumes a lot more space than unencrypted data. As an example, Figure 14-1 shows the plaintext data as a single sentence, but the ciphertext data is about five times as many characters. The size difference varies with different algorithms, but ciphertext data is larger than plaintext data. This takes more space when the encrypted data is saved on media such as a hard disk and consumes more bandwidth when transferred over a network.



NOTE While these examples are using data classifications of Top Secret, Secret, Confidential, and Unclassified, the same concepts apply no matter what data classification labels an organization uses. Typically, an organization would not encrypt any public data, but would often encrypt any other data that is not public.

Regulatory Requirements

Many jurisdictions have implemented regulatory requirements related to cryptography. Chapter 13 covers some of the laws that mandate protection of privacy information. This includes the protection of personally identifiable information (PII) and protected health information (PHI), though some jurisdictions use different labels than PII and PHI. As a reminder, PII is any information that can personally identify an individual. PHI is similar but it also includes information concerning the health status, provision of health care, or payment of health care for an individual. One of the ways that organizations protect PII and PHI is by encrypting data at rest (stored on media such as a hard drive) and data in motion (transmitted over a network).

The Payment Card Industry Data Security Standard (PCI DSS) applies to any organizations that process credit card payments. Chapter 10 provides an overview of PCI DSS in the context of security audits and lists the 6 control objectives and 12 supporting requirements. The overall goal is to protect credit card data. This includes data at rest and data in motion, and PCI DSS discusses the use of encryption as one of the methods to protect the data.

For data in motion, PCI DSS stresses the use of strong cryptography and security protocols such as Transport Layer Security (TLS), Internet Protocol security (IPsec), and Secure Shell (SSH). PCI DSS version 3.1 specifically discourages the use of SSL and early TLS versions because they are not strong cryptographic protocols.

Participating in Security Awareness and Training

As a security expert in your organization, you may be asked to participate in security awareness and training events to share your knowledge. For example, your organization may require users to encrypt certain data and digitally sign all of their e-mails. However, users may not understand the importance of these actions. By sharing your knowledge, you can help raise the overall security awareness within your organization.

User training isn't always formal. Many times, technicians and administrators provide informal training to end users while troubleshooting problems. Even though this is informal, it can often provide lasting meaning to the users.

Enforcing Integrity with Hashing

Chapter 1 introduced hashing as one of the methods used to enforce integrity. Recall that a hash is simply a number created by executing an algorithm against a file or message. As long as the original data is not changed, the hash will always be the same. Knowing this, you can execute the hashing algorithm on the file or message at one point to create a hash, and then later execute the same hashing algorithm again. If the hashes are the same, you know the original data has not been modified. It has retained integrity.



EXAM TIP Hashing algorithms are used to enforce integrity. A hashing algorithm creates a hash. The hash is calculated at two different times, and if the hashes are the same, you know that the data has not lost integrity.

In contrast, if the hashes are different, you know that the file or message has lost integrity. You don't necessarily know why the file is different. A virus may have infected it, someone may have purposely modified it, or it may have simply lost a bit or two as it traveled over a network. However, you do know it is different and you should treat it with suspicion.

Consider Figure 14-2. Lisa is sending a message (“You will pass the SSCP Exam”) to Homer. Before sending it, she calculates the hash and then sends both the message and the hash. Homer receives both the message and the hash. He can then recalculate the hash on the received message and compare the two hashes.

If the recalculated hash is the same as the sent hash, then the message has not lost integrity. If the recalculated hash is different, Homer knows that the message Lisa sent is not the same as the message that he received.



TIP If Lisa wanted to ensure the confidentiality of her message, she could have encrypted it to prevent unauthorized disclosure. However, because this discussion is focused on integrity, the message is not encrypted.

Hashing Algorithms Provide One-Way Encryption

An important point about hashing algorithms is that they provide one-way encryption of data. They create a hash by reducing the message or file to a message digest, which is

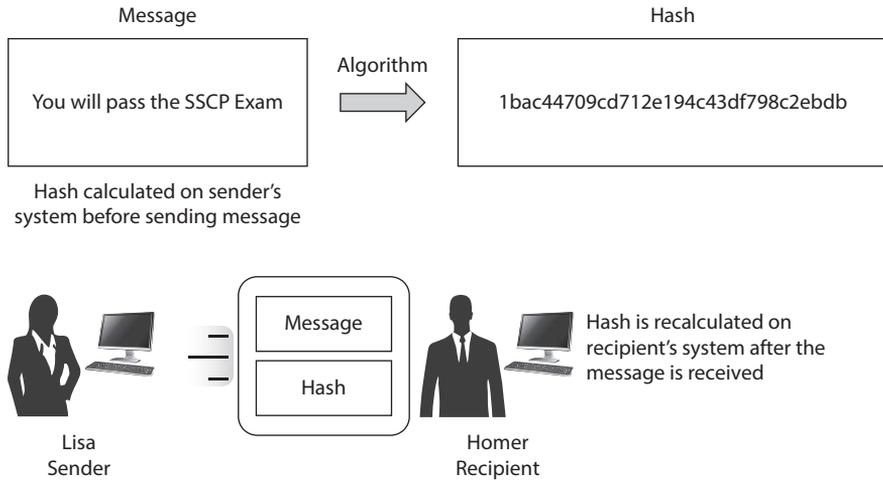


Figure 14-2 Using a hash to verify integrity

a fixed-length number. However, in most cases, you can't reverse this process to discover the original data.

For example, the message “You will pass the SSCP Exam” results in a hash of 1bac44709cd712e194c43df798c2ebdb. However, you can't use this hash to determine the original message. The hash could have come from a 10KB message, a 5MB file, or data of any other size and content. As an example, consider this hash: df076ba5c4ac4d-04cb286bde900c68f2. Just looking at the hash, there's no way that you could determine that it is a hash of a computer file on my computer.



EXAM TIP Hashing algorithms provide one-way encryption and are sometimes referred to as one-way functions. You can create a hash from a file or a message. However, you can't discover the original file or message if you only have the hash.

Hashing Algorithms

Some popular hashing algorithms are Message Digest 5 (MD5), Secure Hashing Algorithm 1 (SHA-1), SHA-2, SHA-3, and Hash-based Message Authentication Code (HMAC).

- MD5** This creates a 128-bit hash, typically displayed as 32 hexadecimal numbers. It is widely used in applications, although security experts have identified vulnerabilities with MD5. The U.S. government considers it cryptographically broken and it is not recommended for use in government applications. The hash shown in Figure 14-2 is the actual MD5 hash of the statement “You will pass the SSCP Exam” created from an online MD5 generator.

- **SHA-1** This is a 160-bit hash. It was commonly used by U.S. government agencies such as the U.S. military until 2010. Due to vulnerabilities, it is no longer recommended for use.
- **SHA-2** This is an improvement over SHA-1 and can create hashes of 224, 256, 384, or 512 bits. Although SHA-2 hasn't been broken, it functions similarly to SHA-1, and many experts believe it's just a matter of time before it is cracked.
- **SHA-3** The U.S. National Institute of Standards and Technology (NIST) sponsored the NIST hash function competition from 2007 to 2012. NIST selected a subset of the Keccak algorithm as SHA-3. The algorithm is completely different from SHA-1 and SHA-2 and creates hashes of 224, 256, 384, or 512 bits.
- **HMAC** The HMAC algorithm works with other hashing algorithms such as MD5 and SHA versions. It uses a shared secret key that adds some randomness to the hash. Only the sender and the receiver know the shared secret key. A primary benefit of HMAC over other hashing protocols is that it adds authenticity.

While MD5 and the SHA algorithms provide integrity, HMAC provides both integrity and authenticity. Figure 14-3 shows the extra step used with HMAC-MD5. In this example, both Lisa and Homer know the secret key is “You Can Pass.” Lisa’s system first calculates the MD5 hash on the message and then uses the secret key to create the HMAC-MD5 hash. She then sends both the message and the hash. Homer’s system then calculates the MD5 hash and uses the secret key to calculate the HMAC-MD5 hash. As long as both hashes are the same, Homer knows that the message has not lost integrity and that Lisa sent it.

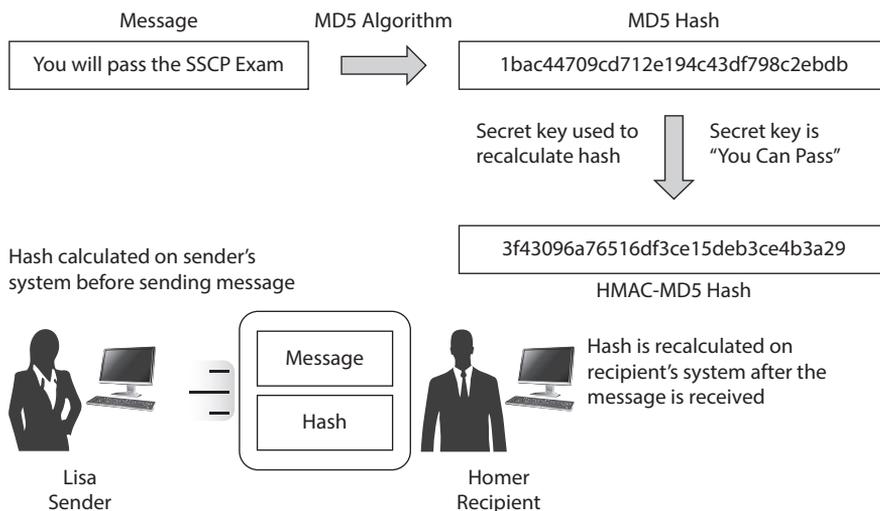


Figure 14-3 Using HMAC to provide integrity and authenticity

In contrast, imagine Bart wanted to impersonate Lisa and send a message to Homer. Because Bart doesn't know the secret key, he might guess one (such as "I won't write on chalkboards"). However, this secret key creates a completely different HMAC-MD5 hash. Homer's system would calculate the HMAC-MD5 hash using the secret key of "You Can Pass." After comparing both, his system would recognize that the keys are different and show an error. This indicates to Homer that Lisa did not send the message.



EXAM TIP MD5 creates a fixed-size 128-bit number (represented as 32 hexadecimal characters) from any message or file. SHA-1 creates 160-bit hashes. SHA-2 and SHA-3 create 224-, 256-, 384-, or 512-bit hashes. HMAC adds authenticity with a shared secret key.

Verifying a Hash

Hashes can be verified automatically or manually. For example, applications that use digital signatures will automatically check the hash to verify the integrity of a message and report the results to the user. The section "Digital Signatures" later in this chapter shows the process used to check the hash automatically.

Other times, you might want to verify the hash of a file manually. For example, if you download a file from a website, you may want to verify that the file you downloaded is the same as the file that the website developers posted. Website developers sometimes calculate the hash of the file before posting it. They then post both the hash and the file. After you download the file, you can manually calculate the hash to see whether it's the same as the posted hash.

As an example, imagine you download a file named `runthis.exe` from a website. On the website, you see that the SHA-1 hash is posted as `286bde900c68f2df076ba5c4ac4d-04cb27a61ffb`. After you download the file, you can calculate the SHA-1 hash of the file and verify it is the same as the posted hash. If it is, you know the file hasn't lost integrity. If it is different, you know the file you downloaded isn't the same as it was when the web developer posted it.

There are many free hashing tools available on the Internet. I've downloaded the `sha1sum` tool, which you can get a copy of here: <http://gcgapremium.com/data/sha1sum.zip>. Figure 14-4 shows a command prompt window with three commands executed. The first one is `dir`, which shows the files in this folder (or directory). I've copied the downloaded file (`runthis.exe`) and the two files used by the `sha1sum` tool to the SSCP folder. The second command (`sha1sum compute runthis.exe`) calculates the hash on the downloaded file. If I compare this to the file posted on the website, I can see that it is the same. This tells me the file has not lost integrity.

As a reminder, no matter how many times I calculate the hash on the same data, I will always get the same hash. You can see this in the commands labeled 3 and 4 in Figure 14-4.

If the hash of the downloaded file is different from the posted hash, you would know that the file you downloaded has lost integrity. You don't necessarily know what caused the problem, but you do know that you shouldn't trust this file. It may have just lost a bit or two during the download, but it also might be infected with malware.

```

C:\SSCP>dir
Volume in drive C is Windows
Volume Serial Number is 260A-97DE

Directory of C:\SSCP
05/12/2015 09:23 AM <DIR> .
05/12/2015 09:23 AM <DIR> ..
02/22/2015 01:57 PM 11,280,580 runthis.exe
05/12/2011 07:47 AM 5,632 SHA1Sum.exe
05/12/2011 07:47 AM 4,608 SHA1Sum_library.dll
3 File(s) 11,290,820 bytes
2 Dir(s) 2,746,597,146,624 bytes free

C:\SSCP>sha1sum compute runthis.exe
286bde900c68f2df076ba5c4ac4d04cb27a61ffb runthis.exe

C:\SSCP>sha1sum compute runthis.exe
286bde900c68f2df076ba5c4ac4d04cb27a61ffb runthis.exe

C:\SSCP>sha1sum compute runthis.exe
286bde900c68f2df076ba5c4ac4d04cb27a61ffb runthis.exe

C:\SSCP>

```

Figure 14-4 Calculating the hash with sha1sum.exe



TIP Md5sum.exe is another freeware tool that calculates hashes. While sha1sum.exe calculates a SHA-1 hash, md5sum.exe calculates the MD5 hash for files. Both are available as free downloads from multiple sites on the Internet.

Salting Passwords

Chapter 5 discusses rainbow table attacks and how salting passwords helps prevent rainbow table attacks. Salting adds a random number of bits to a password. To get a better picture of this, it's valuable to understand how passwords are stored.

Passwords are rarely stored on a system in cleartext. A common way that systems and applications store passwords is in a hashed format. For example, each of the following bullets shows a password and the corresponding SHA-1 hash of the password. Instead of storing the passwords, systems store the hashes.

- **123456** 7c4a8d09ca3762af61e59520943dc26494f8941b
- **IW1llP@33SSCP** f281f938734ccbb10aaf2a17f6ced2198be5582d
- **YouW1llP@33SSCP** c33d3e769a77776bd2557029d0e9a1d6e0db4a28



TIP I mentioned earlier that you normally cannot reverse a hash to discover the original data. The exception is with weak passwords. Without salting, it is trivial for an attacker to discover the actual password from the SHA-1 hash of 7c4a8d09ca3762af61e59520943dc26494f8941b.

When a user logs on and enters a password, the system hashes the password and compares the hash of the entered password with the hash of the stored password. If they're the

same, the system knows the user entered the correct password. This is also useful when transmitting the user's credentials over a network. Instead of sending the user's password, a system can send the hash of the password that the user sent.

While a hash of `7c4a8d09ca3762af61e59520943dc26494f8941b` looks secure (certainly more secure than the weak password of `123456`), attackers found that they could reverse-engineer the hash to discover the password. As mentioned in Chapter 5, they developed rainbow tables, which are large tables of passwords (created from predefined character sets), and calculated hashes of the passwords. Attackers can then perform a simple rainbow table lookup process that compares the hash of a captured password with each of the hashes in the rainbow table. When they have a match, they have discovered the password. You can see that just a few bits changes the hash significantly.

A salt is a set of random bits added to a password before hashing it. These additional bits significantly change the password. For example, by adding the hexadecimal number `9` (`1001` in binary), it changes the hash of the previous three passwords as shown in the following bullets. This makes it difficult for attackers to discover even the weak password of `123456`.

- **123456** `99e2b89f3cda731add47da0c4a7698c1e4c80ba3`
- **IW11IP@33SSCP** `52d02acb0467a856f56c2974cb0d1d5d675042e7`
- **YouW11IP@33SSCP** `949245c0ab1c894eaa8ac42263c8747070707acb`

Exploring Symmetric Encryption

Symmetric encryption uses a single key to encrypt and decrypt data. Symmetric encryption algorithms are constant, but the keys used to encrypt different sets of data are changed frequently to avoid compromising security. For example, the Advanced Encryption Standard (AES) always uses the same algorithm to encrypt and decrypt data, no matter how many times it's used. The key used to encrypt a set of data is also used to decrypt that data.

However, symmetric encryption algorithms use different keys to encrypt and decrypt other sets of data. For example, the algorithm might use a key of `123456` to encrypt and decrypt one block of data, but then use a key of `789321` to encrypt and decrypt another block of data. These keys are shortened for brevity but would be much longer. For example, AES has a default key length of 128 bits.



EXAM TIP Symmetric encryption uses a single key to encrypt and decrypt data. This key should be protected and changed regularly. In contrast, asymmetric encryption uses two keys (called public and private keys), where one key encrypts the information and the other key decrypts the information.

Many applications use symmetric encryption. For example, Chapter 2 presents Kerberos as a network authentication protocol and discusses how it uses tickets for authentication. Kerberos uses symmetric encryption to protect the confidentiality of these tickets.

ROT13

The simplest way to demonstrate symmetric encryption is with ROT13, an encryption scheme used by ancient Romans. I've seen it used occasionally online by people attempting to make their messages harder to read. For example, if someone posts movie spoilers on online forums using ROT13, a reader can't easily read them.

ROT13 rotates letters 13 spaces. If you wanted to post a message saying, "The butler did it," with ROT13, you would rotate each letter forward 13 spaces. If you get to Z, then you continue with A. T rotated 13 spaces is G, h rotated 13 spaces is u, and so on. The ultimate message is "Gur ohgyre qvq vg." Anyone receiving this ciphertext can rotate the letters backward 13 spaces to read the plaintext.



TIP On sites such as <http://rot13.com>, when you enter a message, the site provides the ciphertext in ROT13 form. Similarly, you can enter the ROT13 form and decipher the message.

There are variations on ROT13. For example, you could use a rotation algorithm to encrypt and decrypt the data, but instead of always using the key of 13, you could use a different key at different times. On Monday, for instance, you might use a key of 3 to rotate the letters three spaces forward. To decrypt the data, you would use the same key of 3 but rotate three spaces backward. This allows you to modify the key regularly so that if other people discover the data, they won't necessarily know what the key is and may have trouble decoding it. Although this is a very weak encryption algorithm compared to what's in use today, it does allow you to see easily the difference between the algorithm and the key.

This symmetric rotation algorithm could be defined as follows:

- **Encryption** Rotate X spaces forward to encrypt, where X is the key.
- **Decryption** Rotate X spaces backward to decrypt, where X is the key.
- **Key** The key is any number. Data encrypted with a key of X can be decrypted with the same key of X . In other words, if you use a key of 3 to encrypt the data, you would use the same key of 3 to decrypt the data. The challenge is to ensure that all authorized entities know what key to use.

Encryption algorithms are usually published and well known. For example, if you wanted to look up the encryption algorithm used for any of the symmetric encryptions discussed in this chapter, you could probably find them on the Internet. However, because data is encrypted with different keys, attackers will not be able to read ciphertext with only the algorithm. They also need the key.

Composing and Rotating Keys

Strong encryption algorithms create longer keys and change them often. This prevents attackers from cracking the encryption by guessing the key and ensures that the algorithm can use different keys. This is difficult to conceptualize with large keys, but consider keys of 2 bits, 3 bits, and 4 bits as shown in Table 14-1.

Table 14-1

Number of Bits
Used in Keys

Bits	Calculation	Keys
2	2^2	4
3	2^3	8
4	2^4	16
5	2^5	32

If an algorithm used only 2 bits, it could only create four keys, and attackers could easily crack the code in a brute-force attack using each of the four keys. However, with more bits, the algorithm can create more keys. Notice that adding a bit doubles the number of keys you can create. This is consistent with each bit you add.

Some older encryption algorithms used 56-bit keys, allowing them to create as many as 72,057,594,037,927,936 keys (over 72 quadrillion keys). While the number looks large, it is quite small by today's cryptographic standards. Attackers using basic computers can use a brute-force attack to guess the correct key and crack a code in less than ten minutes. Of course, faster computers can go through the keys quicker.

Current encryption algorithms use larger keys such as 128 bits or 256 bits. For comparison, the same computer that can crack the 56-bit key in less than ten minutes would take about 36 quadrillion years to crack a 128-bit key in a brute-force attack. The trade-off is that larger key sizes require more time to encrypt and decrypt. This won't be noticeable for small sets of data, but will become noticeable when encrypting larger and larger files.

In addition to using longer encryption keys, the encryption algorithm should create random keys and rotate between different keys. Ideally, the algorithm wouldn't ever use the same encryption key twice within the same set of data. If the algorithm did use the same key twice while encrypting a set of data, the code would be easier for an attacker to crack.



TIP Attackers can crack a code when weak keys are used or not changed often. Strong encryption algorithms use large keys and change the keys often.

Comparing Block and Stream Ciphers

Two common types of ciphers used in symmetric encryption are block ciphers and stream ciphers. Both ciphers use a single key to encrypt and decrypt specific pieces of data, but they divide the data differently.

A block cipher encrypts fixed-length groups of bits, or blocks of data, individually. It starts by dividing the plaintext data into fixed-length sizes, such as 128-bit blocks of text. It then encrypts each block individually.

In contrast, a stream cipher encrypts the bits as a stream of data rather than as fixed-length blocks. Stream ciphers are susceptible to more vulnerabilities when not implemented properly. For example, an important rule related to stream ciphers is never to reuse a key. Wired Equivalent Privacy (WEP) is an early wireless security protocol that broke this rule, making it easy to crack. It's worth stressing that stream ciphers are secure when used properly.

Advanced Encryption Standard

Advanced Encryption Standard (AES) is one of the most popular symmetric encryption algorithms. It is a fast, efficient, 128-bit block cipher. NIST spent five years examining multiple encryption algorithms looking for a replacement for an older encryption algorithm. NIST ultimately selected the Rijndael algorithm, and both public and private sectors have adopted it for use. This is similar to how NIST selected a subset of the Keccak algorithm as SHA-3 to replace SHA-1 and SHA-2.

Many different applications use AES for encryption. For example, Wi-Fi Protected Access 2 (WPA2) uses AES to encrypt wireless transmissions. AES is used to encrypt entire hard drives with Trusted Platform Modules (TPMs) and can encrypt data on USB flash drives. These are just a few examples of AES's popularity as a symmetric encryption standard, and it is important to realize why—AES is strong and efficient.



EXAM TIP AES is a strong, efficient symmetric algorithm and is used in many applications. NIST selected it as a replacement for Data Encryption Standard (DES) over 15 other competing algorithms.

AES uses keys of varying sizes to add strength to the cipher:

- AES is sometimes listed as AES-128 and uses 128-bit keys.
- AES-192 uses a 192-bit key.
- AES-256 uses a 256-bit key.

Larger key sizes make the encryption significantly harder to break, but there is an associated resource cost. For example, AES-256 is stronger than AES-128, but systems can encrypt and decrypt AES using 128-bit keys quicker than they can encrypt and decrypt AES-256.



TIP You can't accurately compare the strength of different encryption algorithms based simply on their key sizes alone. For example, imagine if I created a simplistic cipher called Gibson's Cipher and used 2,048-bit keys. It wouldn't be accurate to say this cipher is stronger than AES-256 simply because Gibson's Cipher uses bigger keys. However, it is accurate to say that AES-256 is stronger than AES-128 because it's the same algorithm with larger keys.

Other Symmetric Encryption Algorithms

The following list shows many of the other symmetric encryption algorithms that you may come across. However, none of these have the combined strength and efficiency of AES, so their usage is significantly less common today than it has been in years past. Additionally, attackers and security experts have discovered vulnerabilities in some protocols, so experts recommend discontinuing their use. This includes DES and RC4.

- **DES** DES is an older 64-bit block cipher that uses 56-bit keys. It has been broken and is rarely used anymore. It was a standard selected by the National Bureau of Standards (now known as NIST), but was replaced by AES.
- **3DES** Triple DES (or three DES) improves DES by encrypting data in three passes with three separate keys. It is also a 64-bit block cipher and it was one of the standards evaluated by NIST with AES but was not selected. It supports 56-bit, 112-bit, or 168-bit keys. It is a slower and processor-intensive block cipher, but is strong.
- **Blowfish** Blowfish is a 64-bit block cipher that is still used in many applications. It supports 32-bit to 448-bit keys and was another standard evaluated by NIST as a replacement for DES. It is a strong cipher and is still in use, although its use is being overshadowed by AES.
- **Bcrypt** Many Unix and Linux distributions use bcrypt to protect passwords stored in the shadow password file. It salts the password with additional bits before encrypting it with Blowfish. Bcrypt is one of many key-stretching techniques. Key stretching transforms a potentially weak password into a more secure password.
- **Password-Based Key Derivation Function 2 (PBKDF2)** PBKDF2 is another key-stretching technique. It adds a salt of at least 64 bits and then applies a cryptographic function (such as a cipher or hashing algorithm) multiple times. PBKDF2 can be used for passwords similar to bcrypt or to create stronger cryptographic keys used for symmetric encryption.
- **IDEA** The International Data Encryption Algorithm (IDEA) is a 64-bit block cipher using 128-bit keys that was also submitted as a replacement for DES. Many implementations of Pretty Good Privacy (PGP) used IDEA for encrypting e-mails, but IDEA is rarely used with current implementations of PGP.
- **RC4** Ron Rivest invented Rivest's Cipher (RC4), and it has been a widely used symmetric stream cipher. It has been used with Transport Layer Security (TLS) to encrypt HyperText Transfer Protocol Secure (HTTPS) traffic. Many experts have speculated that RC4 has been cracked, and it is no longer recommended for use.

Table 14-2 summarizes key points about some popular symmetric encryption algorithms.

Algorithm	Block or Stream Cipher	Keys	
Symmetric Encryption Algorithms	AES	128-bit block	128-bit, 192-bit, or 256-bit keys
	DES	64-bit block	56-bit key
	3DES	64-bit block	56-bit, 112-bit, or 168-bit keys
	Blowfish	64-bit block	32-bit to 448-bit keys
	IDEA	64-bit block	128-bit keys
	RC4	Stream cipher	40-bit to 2,048-bit keys

RC4 Weaknesses

While RC4 is widely used, it does have some cryptographic weaknesses. Due to these weaknesses, many experts have been recommending replacing it.

Several reports have indicated that the U.S. National Security Agency (NSA) and the UK Government Communications Headquarters (GCHQ) can crack Secure Sockets Layer (SSL) sessions that use RC4. Most websites have disabled SSL at this point and replaced it with TLS. However, some experts believe that the NSA can crack TLS sessions that use RC4.

RFC 7465 (a proposed standard at this writing) prohibits the use of RC4 with any TLS version due to a variety of cryptographic weaknesses. The RFC states that RC4 is “on the verge of becoming practically exploitable” and “RC4 can no longer be seen as providing a sufficient level of security for TLS sessions.”

TLS can work with other symmetric algorithms, such as 3DES and AES. Some experts and organizations such as Mozilla have recommended using 3DES instead of RC4. While 3DES is slower than RC4, it is more secure. AES would be a better alternative, but some older browsers (such as Internet Explorer 7 and 8) do not support AES.

Microsoft released “Security Advisory 2868725: Recommendation to disable RC4” in 2013. Similarly, Microsoft Knowledge Base article 2868725 recommends disabling RC4 and provides the steps to do so.

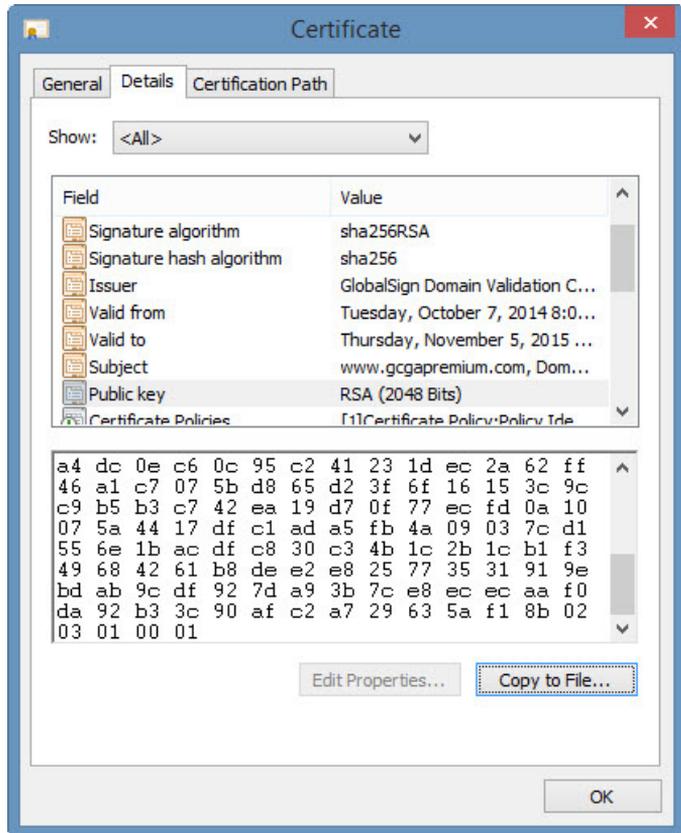
Exploring Asymmetric Encryption

Asymmetric encryption (also known as public key cryptography) uses two matched keys identified as a public key and a private key. There are several important points to remember related to these keys:

- Anything encrypted with the public key can only be decrypted with the matching private key.
- Anything encrypted with the private key can only be decrypted with the matching public key.
- The private key is always kept private and never shared.
- The public key is freely shared and publicly available.

Asymmetric encryption requires a public key infrastructure (PKI, described in more detail later in this chapter) to create, manage, distribute, validate, and revoke certificates. A public key is embedded in a freely shared certificate. For example, Figure 14-5 shows a certificate provided by www.gcgapremium.com after a user requests a secure session. This certificate is using a 2,048-bit public key, and the matching private key is stored privately on the www.gcgapremium.com web server.

Figure 14-5
Certificate with
public key
selected



A common question people ask is, “Which key encrypts and which key decrypts?” The answer depends on the goal. The following are two common uses of these keys in asymmetric encryption:

- Privately share a symmetric key between parties over a public network**
 In this case, the public key encrypts and the private key decrypts. The “Transport Layer Security” section later in this chapter describes how this works.
- Digitally sign an e-mail** In this case, the private key encrypts and the public key decrypts. The “Digital Signatures” section describes how this works.



EXAM TIP Public keys are embedded in certificates and freely shared. The public key is matched with a private key, and the two matched keys work together to encrypt and decrypt data.

RSA

RSA is widely used in public key cryptography for privately sharing symmetric keys (key exchange) and digital signatures (described later in this chapter). Ron Rivest, Adi Shamir, and Leonard Adleman publically described the RSA algorithm in 1978, and its name is derived from the first letters of their last names. In the field of cryptography, it is rather amazing that an algorithm created in 1978 is still in use today.

The reason why it has lasted so long is that it can adapt to using larger keys as computing power increases. Early implementations used smaller keys such as 512 bits. These were secure for a very long time, but can be cracked relatively quickly today. NIST recommends using 1,024-bit keys for any data that needs to be protected through the year 2015. Beyond 2015, NIST recommends using 2,048-bit keys. These recommendations are based on computing power and the likelihood that cryptanalysis techniques will be able to crack 1,024-bit keys at some point in the future.

RSA uses large prime numbers to create secure, matching public and private keys. These large prime numbers are multiplied together to create a composite number that is then manipulated to create the public and private keys. Although it's easy to multiply the two large prime numbers together to create the composite number, it is not easy to factor the composite number to discover the original two prime numbers.

In other words, when the prime numbers are large enough, it's not possible to use the composite number to identify them, as least not in a reasonable amount of time. The original prime numbers remain secret, and more importantly, the private key derived from these numbers remains secret. Due to the increasing speed of computers, the goal of "large enough" prime numbers has been a moving target. Similarly, the size of public and private keys has been steadily increasing to stay ahead.



EXAM TIP RSA uses large prime numbers. Its strength lies in the fact that factoring the composite number created from these prime numbers is computationally infeasible as long as the prime numbers are sufficiently large.

Transport Layer Security

Transport Layer Security (TLS) encrypts Internet transmissions as HTTPS over port 443. TLS is used for other purposes, such as securing virtual private network (VPN) transmissions, but this section is focused on TLS's use within HTTPS. It describes how TLS uses asymmetric encryption to share a symmetric key privately.

It's important to realize that TLS uses both symmetric and asymmetric encryption, as follows:

- **TLS uses asymmetric encryption to share the session key privately** The public and private asymmetric keys are used to encrypt and decrypt the session key.
- **TLS uses symmetric encryption to encrypt session data** The session key is then used to encrypt and decrypt the data within the session.

Figure 14-6 shows the overall process. In this example, Bart is establishing a secure session with a web-based server, such as a web-based e-mail site. Bart establishes an HTTPS

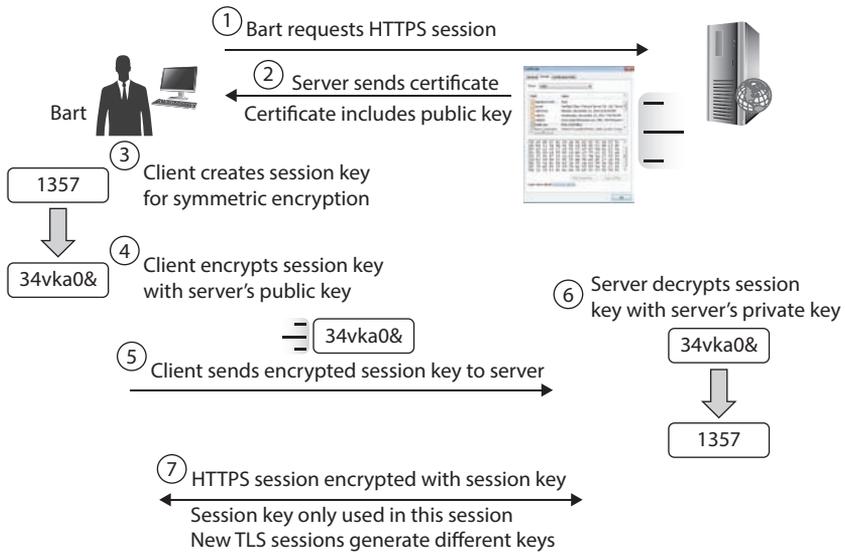


Figure 14-6 Creating a secure HTTPS session with TLS

session before entering his credentials to access the e-mail, and the session ensures his credentials aren't sent over the Internet in cleartext. As you look at the figure, you can trace the following steps:

1. Bart initiates the process by requesting an HTTPS session. For example, he could simply click a link on a website page, such as an account logon page.
2. The server responds by sending its certificate with the server's public key embedded within it. This public key is matched with a private key known only by the server.
3. After receiving the certificate, the client (Bart's computer) creates a session key (which will be used as a symmetric key) that will be used later to encrypt session data. At this point, only the client knows the session key. For brevity, this discussion is using a simplistic session key of 1357, but it would be much larger in practice. For example, TLS can use a 256-bit AES key.



TIP Although Figure 14-6 doesn't show it, the client validates the certificate before using it. This includes ensuring that the certificate hasn't expired and that the certificate authority hasn't revoked the certificate. The "Validating Certificates" section later in this chapter describes this process in a little more depth.

4. The client encrypts the session key with the server's public key. Recall that only the matching private key (which only the server knows) can decrypt the encrypted session key.

5. The client sends the encrypted session key to the server over the Internet. Even if an attacker captures this encrypted session key, it won't be of any use because the attacker won't have the matching private key needed to decrypt the session key.
6. When the server receives the encrypted session key, it decrypts it with its private key. At this point, both the client and the server know the session key, but no one else knows it.
7. The session key encrypts data within the HTTPS session using symmetric encryption. If attackers intercept any of these pages, they won't be able to read the data because they don't have the session key. After the session ends, the client and server discard the session key. If the user initiates another HTTPS session, the two systems repeat the process. Session keys have the shortest lifespan when compared to other keys.

A logical question to ask here is why not use asymmetric encryption to encrypt the entire session. The reason is that symmetric encryption is much faster than asymmetric encryption. Depending on the algorithm used, symmetric encryption can be as much as 100 times faster than asymmetric encryption.

Secure Sockets Layer

Secure Sockets Layer (SSL) is the original protocol used to encrypt HTTPS sessions. Netscape published it in 1996, but there wasn't any formal method of updating it. The Internet Engineering Task Force (IETF) ultimately created TLS as the replacement and has been steadily updating and improving TLS over the years.



TIP SSL has not been defined as a formal RFC standard. However, the IETF released RFC 6101 as a Historic document that describes SSL in depth.

A Google security team disclosed a vulnerability with SSL named POODLE (Padding Oracle On Downgraded Legacy Encryption) in late 2014. While many organizations had already switched to TLS by this time, this announcement prompted more to do so. Soon afterward, many organizations disabled SSL in their applications, and many applications now specifically block its use.

Diffie-Hellman

Diffie-Hellman is one of the first public key algorithms used to privately share symmetric keys between two systems. Whitfield Diffie and Martin Hellman published the method in 1976, and it has been adopted in many public key algorithms. For example, TLS supports the use of Diffie-Hellman to exchange session keys.

The focus of Diffie-Hellman is on secure key exchange. As a comparison, RSA provides secure key exchange and is also a digital signature algorithm used to create digital signatures. The "Digital Signatures" section later in this chapter describes digital signatures.

Elliptic Curve Cryptography

Elliptic curve cryptography (ECC) uses a mathematically defined curve to create public and private keys. A benefit of ECC is that it provides strong security with smaller key sizes when compared to other methods such as RSA. ECC is often used in smaller mobile devices because the smaller keys require less processing power.

Several algorithms support ECC. For example, the Elliptic Curve Digital Signature Algorithm (ECDSA) uses ECC methods for digital signatures. Similarly, elliptic curve Diffie-Hellman (ECDH) is a variant of Diffie-Hellman that uses ECC for key exchange.



NOTE Reports in 2013 indicated that the NSA influenced NIST to include a flaw in a recommended elliptic curve. Later that year, RSA Security LLC issued an advisory recommending the discontinued use of ECC based on Dual Elliptic Curve Deterministic Random Bit Generation (Dual_EC_DRBG).

Secure Shell

Secure Shell (SSH) uses public key cryptography to encrypt data in motion. Chapter 3 discusses the use of SSH in more depth. It is a secure alternative to Telnet and similar protocols such as rlogin, rsh, and rexec, which all send credentials over a network in cleartext. Secure File Transfer Protocol (SFTP) and Secure Copy (SCP) both use SSH to encrypt the data transfers.

Protecting E-mail with S/MIME

Secure/Multipurpose Internet Mail Extensions (S/MIME) is the standard used to encrypt and digitally sign e-mail. Almost all e-mail applications that support encryption and digital signatures use S/MIME. The following sections describe how S/MIME uses digital signatures and encrypts e-mail.

Nonrepudiation

Nonrepudiation prevents individuals from believably denying they took an action. This is similar to how you cannot believably deny making a charge on your credit card if you provided your signature during the transaction. Of course, you can deny it, but your claim isn't believable.

Electronic signatures provide another example of nonrepudiation. Mortgage companies, banks, realtors, and many other organizations use electronic signatures when sharing documents over the Internet. Electronic signatures are valid and legally binding, just as a signature made with a pen on paper is legally binding.

DocuSign is one of the companies many organizations use for online electronic signatures. Here's how it works. A company representative uploads documents to a DocuSign site and identifies the customer that needs to sign it, using the customer's e-mail address. The representative then marks the document where the customer needs to initial or sign it. The customer receives an e-mail with a link for the document. They agree to use the electronic signature and follow the instructions to sign and initial the document. The

customer cannot later say they never received the document, because the electronic signature provides legal proof that they did.

Digital signatures (described in the next section) are another method used for nonrepudiation.

Digital Signatures

A digital signature is an encrypted hash of a message. It uses asymmetric encryption and is encrypted with the sender's private key. Digital signatures provide several security benefits, including authentication, integrity, and nonrepudiation.

- **Authentication** The digital signature provides proof to the recipient that the claimed sender is the actual person who sent the e-mail. In this context, the digital signature provides authenticity.
- **Integrity** The e-mail is hashed before it is sent, and the hash is included with the e-mail. The recipient hashes the received message and verifies that the hashes are the same. If the hashes are the same, that verifies that the message has not been modified. In other words, it verifies the message has not lost integrity.
- **Nonrepudiation** The sender creates the digital signature with the sender's private key. Because only the sender has the sender's private key, the sender cannot deny sending it.



TIP A Digital Signature Algorithm (DSA) creates digital signatures. RSA includes a DSA component, and many applications use RSA to create digital signatures. The Elliptic Curve Digital Signature Algorithm (ECDSA) is another DSA, but unlike RSA, ECDSA does not include a key-exchange component.

Figure 14-7 shows the overall process of creating and using a digital signature. In this example, Maggie is sending an e-mail to Lisa and using the digital signature for authentication, integrity, and nonrepudiation. As you look at the figure, you can trace the following steps:

1. Maggie, the sender, first creates the message. In this example the message is, "A secret to success is setting goals."
2. A hashing algorithm creates a hash from the message.
3. The hash is encrypted with the sender's private key, creating the digital signature. As mentioned in the "Enforcing Integrity with Hashing" section earlier in this chapter, this hash provides integrity.

There are two key points to remember at this point:

- a. Only the sender has the sender's private key.
 - b. Only the sender's public key can decrypt the encrypted hash.
4. Maggie sends both the message and the encrypted hash to Lisa, the recipient.

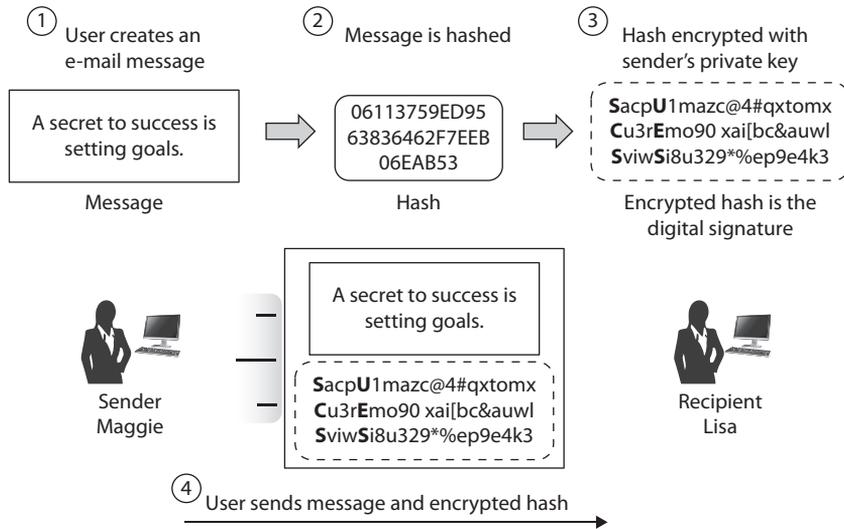


Figure 14-7 Creating a digital signature

Note that the message isn't encrypted as part of the digital signature. In other words, the digital signature does not provide confidentiality. However, if it is important to maintain confidentiality of the message, you can encrypt an e-mail. The "Encrypting E-mail" section later in this chapter shows the encryption process. While it is possible to encrypt and digitally sign the same e-mail message, this section focuses only on the digital signature.

When the recipient receives a digitally signed e-mail, the receiving system verifies the digital signature. Figure 14-8 shows the steps taken after a user receives a digitally signed e-mail. As you look at the figure, you can trace the following steps:

1. The recipient retrieves the sender's public key by retrieving the sender's certificate. Within an internal network (such as in a Microsoft domain), this process is automated without requiring any user interaction. In other situations, users may be required to install the sender's certificate on their system manually. The certificate only needs to be retrieved once.
2. The digital signature is decrypted with the sender's public key. When successful, this step shows the original hash of the message created on the sender's system. This is an important step. It will succeed only if the sender's private key encrypted the hash.
 - a. If the decryption of the digital signature is successful, it proves that it was encrypted with the sender's private key. Because only the sender has the sender's private key, the sender must have encrypted it. This provides both

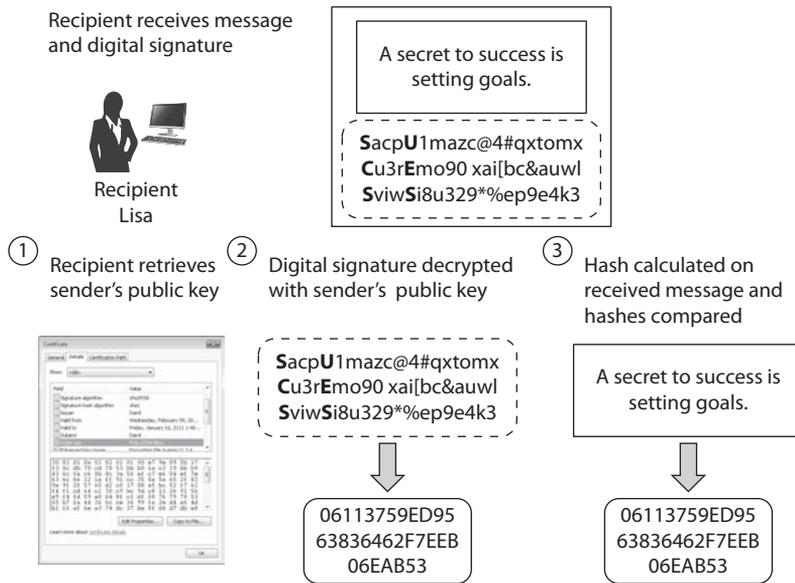


Figure 14-8 Verifying a digital signature

authentication and nonrepudiation. The recipient knows the sender sent the message (authentication), and the sender cannot later deny sending the message (nonrepudiation).

- b. In contrast, if someone tried to impersonate the sender and included a bogus digital signature, this step would fail. The impersonated sender's public key would not decrypt the digital signature, so the process would show an error.
3. The recipient's system calculates the hash on the received message. It then compares the calculated hash with the decrypted hash from the digital signature.
 - a. If both the received hash and the calculated hash are the same, the message has not been modified (or has not lost integrity).
 - b. If the hashes are different, the recipient knows that the received message is not the same as the sent message. The recipient wouldn't know how the message was modified. The differences could be due to tampering from an attacker, a problem with the network, hardware, or software, or something else. However, the recipient knows the content of the message should not be trusted.



EXAM TIP The primary purpose of a digital signature is to provide authentication. However, a digital signature also provides integrity and nonrepudiation.

Encrypting E-mail

You can also use asymmetric encryption to encrypt e-mail. In simplistic terms, here's what happens:

1. The sender retrieves the recipient's public key.
2. The sender encrypts the e-mail with the recipient's public key.
3. The recipient decrypts the e-mail with the recipient's private key.



EXAM TIP An important point to remember about e-mail encryption is that the sender requires the recipient's public key.

Although the preceding steps simplify the process, e-mail applications implement the encryption process with some additional steps. The following explanation provides the underlying details.

In the “Transport Layer Security” section, you saw how TLS uses public key cryptography to privately share a session key and how the session key is used to encrypt the session data. E-mail encryption typically uses a similar process. The primary difference is in how the sender retrieves the certificate holding the recipient's public key.

The certificate retrieval process is automated in many internal networks. For example, in a Microsoft domain, the user could use Microsoft Outlook to draft an e-mail similar to that shown in Figure 14-9 and click the Encrypt button. When the user clicks Send, Microsoft Outlook attempts to retrieve the certificate of the recipient based on the e-mail

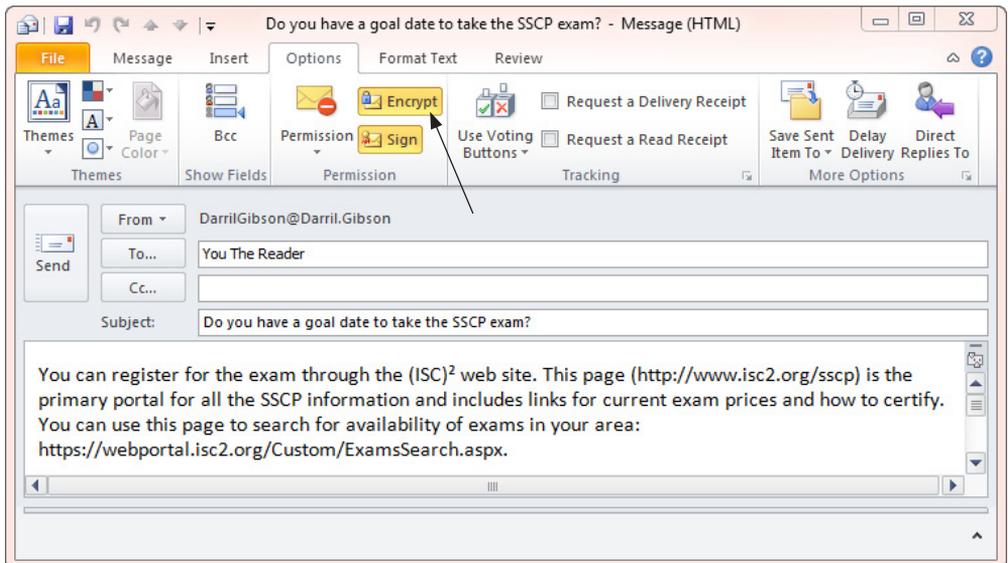


Figure 14-9 Encrypting an e-mail in Microsoft Outlook

address. In other situations, users may be required to install the recipient's certificate on their system manually.



NOTE Figure 14-9 also shows the Sign button, which is used to create a digital signature. Although it's technically possible to encrypt an e-mail without digitally signing it, e-mail applications rarely allow you to do so. Instead, if you try to encrypt an e-mail, the application will force you to add a digital signature.

Once the sender's system has the recipient's certificate, it uses the steps shown in Figure 14-10 and described in the following list to encrypt the e-mail:

1. The sender's system creates a session key. This is similar to how TLS creates a session key. For brevity, this discussion is using a simplistic session key of 2468, but the key would be much larger in practice.
2. The sender's system uses symmetric encryption to encrypt the e-mail with the session key.
3. The sender's system uses asymmetric encryption to encrypt the session key with the recipient's public key.
4. The sender's system sends the encrypted e-mail and the encrypted session key to the recipient.
5. The recipient's system retrieves the recipient's private key.
6. The recipient's system uses the private key to decrypt the encrypted session key.
7. The recipient's system decrypts the e-mail content by using the decrypted session key.

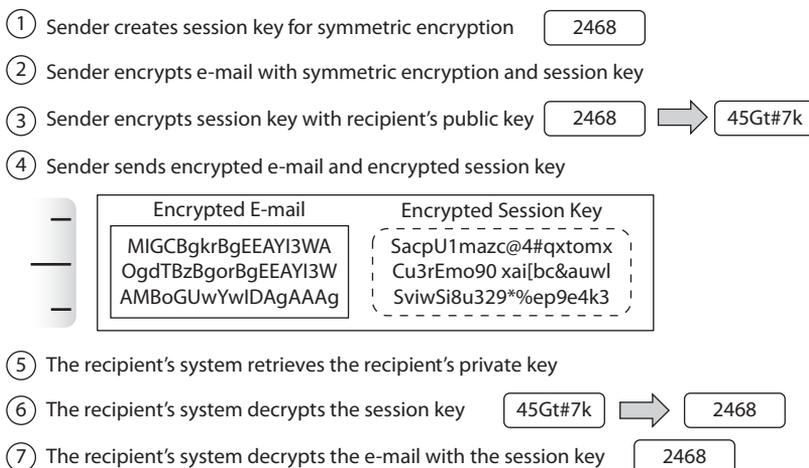


Figure 14-10 Encrypting an e-mail with symmetric and asymmetric encryption

Like the encryption process, the decryption process is automated. For example, the recipient may only need to open the e-mail to view the decrypted content.



TIP Symmetric encryption is significantly faster than asymmetric encryption. You can improve encryption performance by using both symmetric and asymmetric encryption to encrypt the e-mail.

Pretty Good Privacy (PGP)

PGP can encrypt, decrypt, and digitally sign e-mails between individuals. PGP is an application, originally created by Philip Zimmermann, that provides confidentiality through encryption, as well as authentication, with the use of digital signatures.

For example, if you and I wanted to be able to send e-mails between each other and keep them confidential, we could use PGP. I could send you my certificate, including my public key, and you could add it to PGP running on your system. Similarly, you would send me your certificate.

PGP creates a session key, encrypts the e-mail with it, and then uses the public key to encrypt the session key. When my system receives the encrypted e-mail and the encrypted session key from you, PGP uses my matching private key to decrypt the session key and decrypts the e-mail with the decrypted session key. Of course, I could use the same process to send you an encrypted e-mail.

In addition, either of us could digitally sign our e-mails to prove that we sent them. For example, I could use PGP to create a digital signature with my private key. If you had a certificate with my public key, PGP would use it to decrypt the digital signature, providing you with assurances that it was created from my private key.

Other Encryption Schemes

Many encryption schemes fit neatly into the symmetric or asymmetric categories, but a few fall outside the scope of these categories. For example, cryptographers can use steganography to manipulate files and embed detailed messages within them. However, steganography simply hides the data within the file. IPsec is another encryption scheme that doesn't fit neatly into the symmetric and asymmetric categories, but it is important in the context of cryptography.

Steganography

Steganography is the practice of hiding data within data, or hiding data in plain sight. It doesn't actually encrypt the data, but it can provide a level of confidentiality. If someone knows the message is embedded in the data, they can retrieve it.

As an example, some programs modify the least significant bit of individual bytes within a file. For example, consider the byte represented in Table 14-3. The value of the byte is 255 (eight 1's), but if you modify the value of the least significant bit (2^0), the new value is 254. If this value is used to represent the tint of a color or one of the values in a sound file, the change between 255 and 254 is so slight that it is indistinguishable.

2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0	Decimal
1	1	1	1	1	1	1	1	255
1	1	1	1	1	1	1	0	254

Table 14-3 Modifying the Least Significant Bit

If someone modified the last bit of each byte in a 1MB file, he or she could embed a message in these 1 million bits. As an example, the letter A is represented in hexadecimal as 41 and binary as 0100 0001. A person could embed the letter A by ensuring the least significant bit in 8 bytes includes these bits.

Table 14-4 shows the least significant bit from 8 bytes in a graphics file. To embed the letter A in these bytes, a steganography tool ensures that the last bit in these 8 bytes is 0100 0001. Notice that the first byte ends with a 1, so it is changed to a 0. The second byte ends in a 1, so it doesn't need to be changed. The steganography tool would change the last bit in bytes 1, 4, and 7, and leave the other bits the same. This shows how you can embed the letter A in a file. It's possible to repeat this process for a full message.

The changes are not distinguishable to most users. However, when people know what to look for, they will be able to find and read the message. For example, it's possible for someone in one country to post a picture on a website with an embedded message and enable someone in another country to retrieve the picture and read the message.

While most people will overlook the hidden message, security professionals use steganalysis techniques to detect the use of steganography. A common method is the use of hashing. If you have a hash of the original file, you can create a hash from the current version. If the two hashes are different, that indicates the file has been modified and may have a hidden message embedded within it. Additional tools allow experts to analyze the bytes within the file to detect the message.



EXAM TIP Steganography attempts to hide information by embedding it into a file such as a picture or streamed audio or video file. Steganalysis techniques attempt to detect the use of steganography.

	Orig 2^0	Letter A	Modify?	2^0 with Message
Last bit in byte 1	1	0	Yes	0
Last bit in byte 2	1	1	No	1
Last bit in byte 3	0	0	No	0
Last bit in byte 4	1	0	Yes	0
Last bit in byte 5	0	0	No	0
Last bit in byte 6	0	0	No	0
Last bit in byte 7	1	0	Yes	0
Last bit in byte 8	1	1	No	1

Table 14-4 Embedding the Letter A

IPsec

Chapter 3 introduces Internet Protocol security (IPsec) and explains how IPsec provides security for data traveling over a network. The two protocols within IPsec are Authentication Header (AH) and Encapsulating Security Protocol (ESP). AH provides authentication and integrity, while ESP provides confidentiality, authentication, and integrity. You can use AH alone, ESP alone, or both protocols together.

AH provides data origin authentication by assuring both systems of the identity of the other system. It creates a hash derived from packet data and uses this hash to provide integrity in addition to authentication. AH does not encrypt the data or provide confidentiality.

ESP encrypts the data within the packets to provide confidentiality. Although it's possible to configure ESP without integrity, this is not recommended. In other words, when you use ESP for encryption, you will also include the same features provided by AH (authentication and integrity).

Chapter 4 presents IPsec as a tunneling protocol for VPNs. IPsec can be used with Layer 2 Tunneling Protocol (L2TP/IPsec) or by itself to create a secure connection between a VPN client and a VPN server.



TIP Both IPv4 and IPv6 support IPsec. IPv6 mandates the support of IPsec for any IPv6 applications. However, IPsec is an extension to IPv4, so there are compatibility problems in some situations. For example, IPsec traffic can't go through a system using Network Address Translation (NAT).

Public Key Infrastructure

A public key infrastructure (PKI) includes all the components necessary to create, manage, distribute, validate, and revoke certificates. The “Exploring Asymmetric Encryption” section earlier in this chapter explained the use of public and private keys to encrypt information and digitally sign e-mails. The PKI provides the necessary support for these processes, including distributing public keys in digital certificates.

Although you aren't expected to create and manage a PKI for the SSCP exam, you should understand the basic components. This section covers certificates, certificate authorities (CAs), and certificate validation.

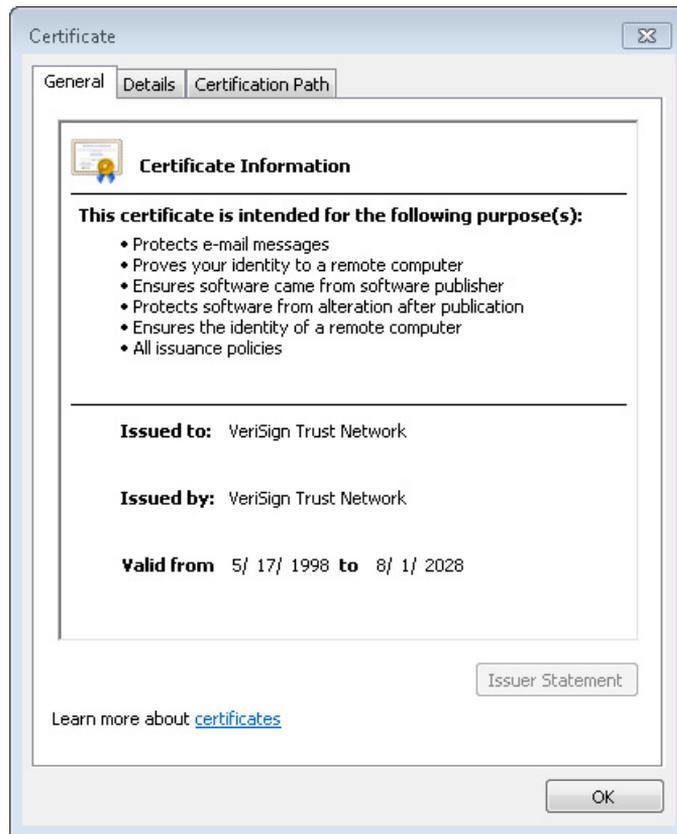


TIP A PKI is created based on the X.509 protocol. The X.509 standard identifies many of the components and formats used by the PKI, CAs, and certificates.

Certificates

The “Exploring Asymmetric Encryption” section earlier in this chapter presented some information on certificates. For example, Figure 14-5 showed a certificate with the public key selected. Although the distribution of the public key is an important purpose of certificates, they do have many other uses. Figure 14-11 shows a certificate from VeriSign with multiple purposes.

Figure 14-11
Certificate
with multiple
purposes



The following are some of the common purposes of certificates:

- **Authentication** Certificates can prove the identity of users and computers. For example, smart cards have embedded certificates used to prove the identity of the user. Remote systems can prove their identity to your system, and your system proves its identity to remote systems with the use of certificates.
- **Encryption** Certificates can be used to assist with encryption. For example, the public and private keys encrypt a symmetric key so that it can be privately shared.
- **Protecting e-mail** This protection includes both encryption and digital signatures.
- **Code signing** Applications and scripts can be digitally signed. This provides authentication of the developer and integrity for the code. If the code is modified (such as with malware), the signature embedded in the certificate will no longer match.



TIP As mentioned in Chapter 5, attackers have been signing malware with code signing certificates, which they stole from others. Extended Validation (EV) Code Signing certificates help mitigate this problem. The Certificate Authority/Browser Forum (CABF) released standards for EV Code Signing certificates in 2012.

When entities want a certificate, they provide the required information to a CA in the form of a request. The CA embeds this information in the certificate.

Figures 14-12, 14-13, and 14-14 show the different tabs of a certificate, giving you an idea of the information that is available. Figure 14-14 shows that GlobalSign issued this certificate to www.gcgapremium.com.

Certificates include the following information:

- Who it was issued to (such as a website, a server, or a user)
- Who issued it (the CA)

Figure 14-12
Viewing the
General tab
of a certificate

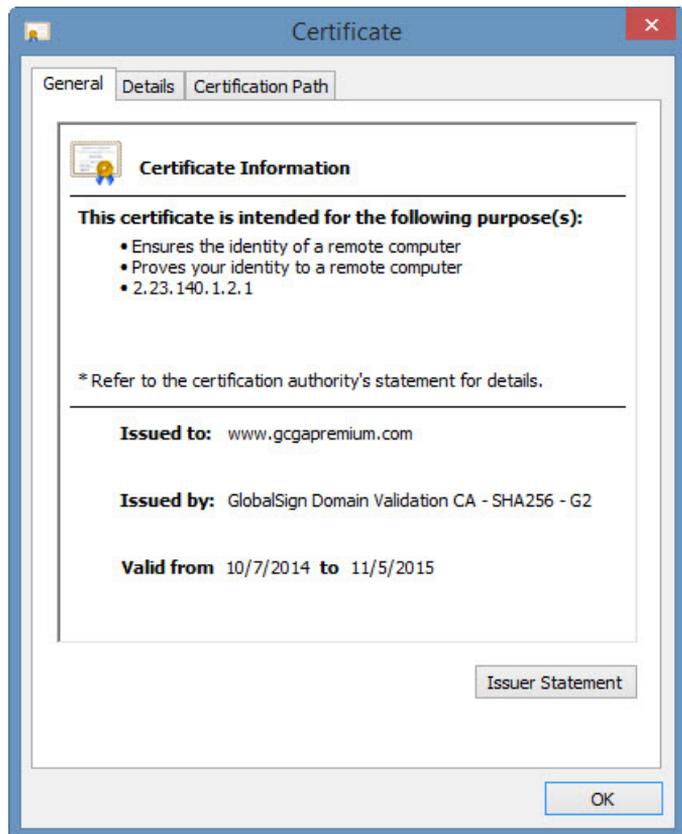
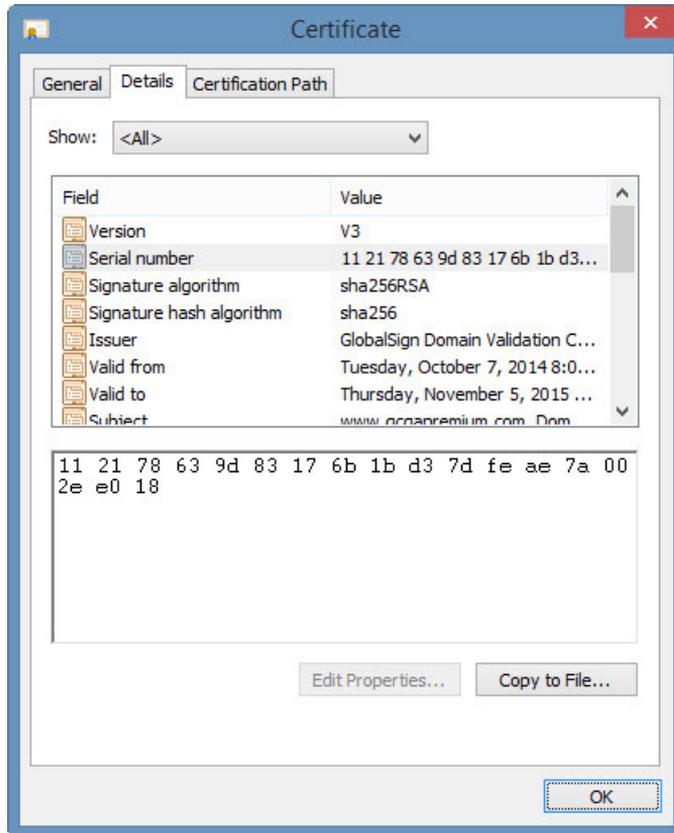


Figure 14-13

Viewing the
Details tab of
a certificate



- Validity dates (which allow the certificate to expire)
- Serial number (used to uniquely identify the certificate)
- Public key (matched to a private key that is not included in the certificate)
- A certification path showing the trust chain of the certificate to the root CA

Certificate Authority

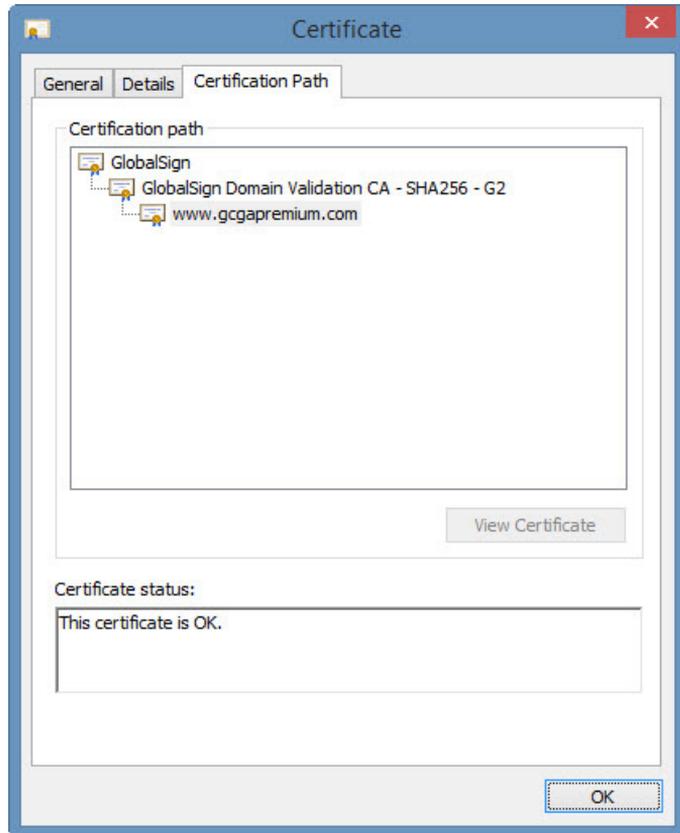
A CA issues and manages certificates through their lifetimes. CAs can be public (such as VeriSign) or private within an organization. Public CAs sell and validate certificates and provide assurances to users that certificates are valid.



TIP A large organization can use a registration authority (RA). An RA accepts certificate requests, verifies the data, and passes the request on to the CA. The RA is optional and never issues certificates.

Figure 14-14

Viewing the Certification Path tab of a certificate



For example, if you wanted to do business on the Internet via your own website, you could purchase a certificate from a public CA to create a secure HTTPS session. Then when a customer wanted to make a purchase, your website would send the certificate to the customer's computer to initiate the secure session.



TIP The "Transport Layer Security" section earlier in this chapter describes the process of establishing an HTTPS session.

First, you would create a matching public key and private key. Many applications are available to create these matching keys. You would then send the public key and other required information to the CA. The CA would create the certificate and send it back to you, and you would then use it to configure your website.

It's also possible to create a private CA and issue certificates internally. The benefit is that the organization with a private CA doesn't need to purchase certificates from a public CA. However, the drawback is that certificates issued by a private CA aren't trusted by default. If external clients receive these certificates, they will receive an error indicating that the certificate is not trusted.

Understanding the Certificate Trust Chain

An important part of the PKI is the trust chain, also known as a hierarchy of trust. For example, if you look back to Figure 14-14, you can see that the Certification Path tab shows two levels. GlobalSign is the top level and in this example is the root CA. The root CA issues certificates to other CAs in the trust chain, known as intermediate CAs.

In Figure 14-14, GlobalSign Domain Validation CA SHA256 - G2 is an intermediate CA. The intermediate CA then issued a certificate to `www.gcgapremium.com`. Systems that trust the root CA or the intermediate CA will automatically trust the certificate issued to `www.gcgapremium.com`.



TIP It's possible to have multiple layers of intermediate CAs. In other words, a root CA can issue a certificate to an intermediate CA. The intermediate CA could then issue certificates to second-level intermediate CAs. This can continue for as many levels as the organization needs, although more than five are rarely necessary. The lowest-level CA issues certificates to others, such as websites or servers.

As long as one of the CAs in the certificate trust chain is trusted, certificates issued by the CAs are also trusted. For example, if the root CA is trusted, certificates issued by any of the CAs in the trust chain are also trusted.

Understanding Trusted Root Certification Authorities

You may be wondering how a CA is trusted. The primary method is by installing certificates from the root CA onto a system's trusted certificate authorities store. For example, Figure 14-15 shows the Trusted Root Certification Authorities store for a Windows 7 system. Certificates in this store are from root CAs, and each certificate represents a certificate trust chain. Any certificate issued by any CA in this trust chain is automatically trusted.

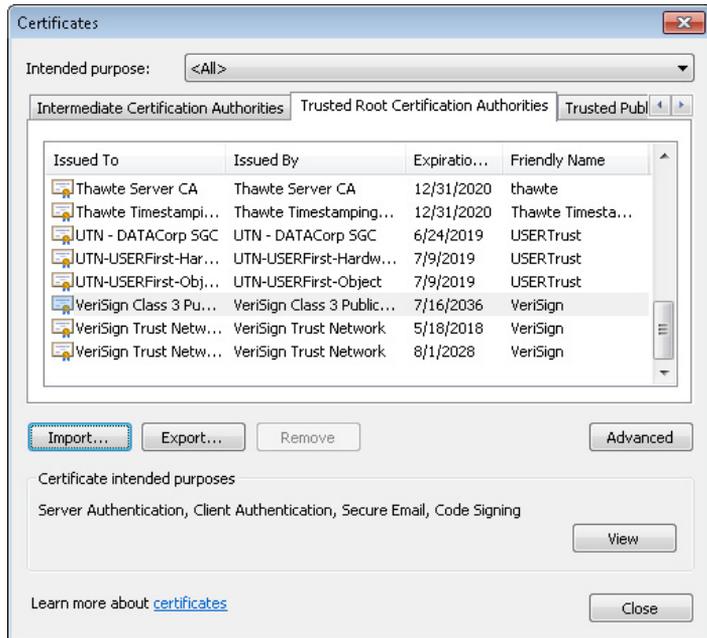
This is similar in concept to a driver's license issued by a Department of Motor Vehicles (DMV). If you need to prove your name or age, you can show your driver's license. Because people trust that the DMV has verified your information before giving you a driver's license, and your driver's license was issued by the DMV, most people trust the validity of your driver's license. On the other hand, if you bought an identification card from Darril's Cheap IDs, many people may question the validity of information on that ID card.

The root CA is similar to the DMV. Any certificates issued by the root CA are trusted because the root CA is trusted, just as licenses issued by the DMV are trusted because the DMV is trusted.



TIP There are other ways you can add certificates to a system so that they are trusted. However, the point here is that certificates will automatically be trusted if they're in a trust chain. E-commerce sites purchase certificates from CAs that are in the trust chain so that their certificates will automatically be trusted by customer computers.

Figure 14-15
Trusted root
certification
authorities store



Revoking Certificates

Most certificates will have a lifetime identified in the validity dates. However, in some situations the certificate is revoked before it expires. For example, the CA will revoke the certificate if the matching private key of a certificate is compromised. The CA maintains a list of all revoked certificates in a certification revocation list (CRL) and publishes the CRL as a V2 certificate. Most other certificates are V3 certificates. The CRL lists the serial numbers of certificates revoked by the CA.

Certificates include a list of CRL distribution points. Clients can request a copy of the CRL through these distribution points and verify that the certificate they are using is not listed on the CRL. This verification is often done as part of the normal process of validating a certificate.

Validating Certificates

When systems receive certificates from a website, they go through several checks to ensure that the certificate is valid. Figure 14-16 shows the overall steps that occur, which are as follows:

1. The user requests the HTTPS session.
2. The server responds with the certificate.
3. The user's system examines the certificate. It ensures that the certificate hasn't expired (based on the validity dates) and ensures that the website that issued the

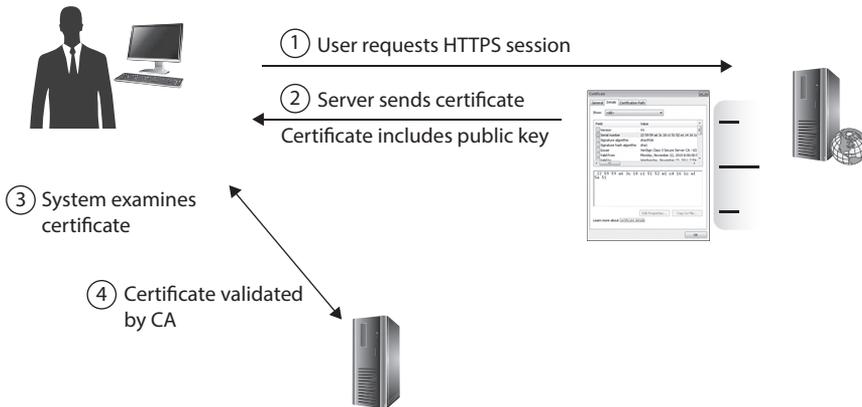


Figure 14-16 Validating a certificate

certificate is the same as the one identified in the certificate. It will also ensure that a CA in the certificate trust chain is trusted.

4. The user's system checks with the CA to ensure that the certificate has not been revoked. A common method is by requesting a copy of the CRL through a CRL distribution point. If the CRL includes the serial number of the certificate provided by the website, the client system knows the CA has revoked the certificate and it shouldn't be trusted.

Another method of checking to see whether the certificate is revoked is with the Online Certificate Status Protocol (OCSP). The client system first identifies the serial number of the certificate received from the website. Then, instead of requesting a copy of the CRL, the client sends the serial number to a server known as an OCSP responder (maintained by the CA). The OCSP responder identifies the health of the certificate with an answer of "good," "revoked," or "unknown." A response of "unknown" might indicate the certificate is a forgery. The benefit of OCSP over a CRL is that the entire CRL doesn't need to be transmitted.

If any of these checks fails, the user will see an error indicating what check failed. The errors are different based on the browser that the user is using.

Key Escrow

Key escrow refers to storing a copy of a private key for safekeeping. It can be retrieved if the original key is lost or unavailable. As an example, if an employee encrypts data and then leaves the organization, the data may not be accessible if the organization hasn't enabled a key escrow method.

Recovery agents are the individuals authorized to recover keys held in escrow. Typically, keys held in escrow are kept within the organization. However, there are some situations where the keys are held in escrow by a third party. If the keys are kept within the

organization, the recovery agent is someone within the organization. If the keys are held by a third party, the recovery agent is typically someone in the third-party organization.

Alternative Certificate Trusts

The “Understanding the Certificate Trust Chain” section described a hierarchical trust chain. However, this isn’t the only type of trust. Two additional trust methods are a web of trust and a cross-certification trust.

A hierarchical trust chain allows CAs to issue certificates and for entities to verify the certificates are valid. For example, imagine a client trusts a root CA. The client will automatically trust certificates from the root CA or any intermediate CAs in the trust chain. Additionally, the client can validate the certificate by either retrieving the CRL or with OSCP.

In contrast, anyone can create self-signed certificates. Because the person who creates a self-signed certificate also signs the certificates, there’s no way to validate them. A web of trust is a decentralized trust model and provides one way to solve this problem.

As a simple example, imagine that Lisa, Homer, Marge, and Bart each create their own self-signed certificates to encrypt e-mails and create digital signatures. They exchange these certificates with each other and ask everyone in their group to trust them. This is a web of trust.

Later, Lisa and Maggie exchange certificates and start exchanging secure e-mails. Lisa chooses to trust Maggie’s certificate and could introduce Maggie’s certificate to the group, encouraging everyone to trust it. If Bart (or anyone else in the web of trust) sends a certificate to Maggie, she might not trust it because Bart is unknown to her. However, because Lisa trusts it, this might be enough for Maggie to trust it.

In the early days of PGP, its biggest challenge was the distribution of certificates. Phil Zimmermann, the creator of PGP, recognized that there wasn’t a reliable method of proving identities. He introduced the web of trust concept in an early manual for PGP. Admittedly, the details are more technical than described in the example with Lisa’s group, but the concept is similar. Designated individuals are authorized to introduce certificates to the web of trust. If enough people trust a certificate, it provides a level of assurance. For example, if six of your friends trust a certificate, you might choose to trust it.



TIP Current implementations of PGP can use a public key infrastructure to retrieve and validate certificates.

Cross-certification allows organizations with a private PKI and private CA to trust entities in another private PKI. As an example, imagine that Globex Corporation has created its own PKI to issue and manage certificates for Globex employees. Similarly, imagine that the Springfield University has created a private PKI for its students, staff, and faculty. By default, these certificates are not trusted outside the organization. If an employee at Globex sends a certificate to someone at Springfield University, it creates errors because the certificate isn’t trusted and cannot be validated.

However, if the two organizations create a cross-certification agreement, it creates a mutual trust relationship. Users in both organizations automatically trust the certificates

from the other organization. A cross-certification agreement will also implement methods allowing systems to validate the certificates.

Comparing Cryptanalysis Attacks

Cryptanalysis is the process of deciphering codes through analysis to compromise the confidentiality or integrity of data. Authorized users of encrypted data have the tools to decrypt and access data when needed. However, unauthorized users need to use cryptanalysis to decrypt and access the data.



EXAM TIP Both attackers and researchers use cryptanalysis techniques. Researchers use cryptanalysis techniques to search for weaknesses in cryptography with the goal of improving cryptographic methods. Attackers search for weaknesses with the goal of exploiting them.

Most cryptographic algorithms are published, allowing anyone to study them. Security experts often perform cryptanalysis on protocols to identify weaknesses based on new data available. As an example, SSL was thought to be secure for many years. However, a security team at Google discovered a significant vulnerability and published a paper on the POODLE attack in late 2014. This prompted many organizations to switch over to TLS sooner than they might have otherwise done. Similarly, the federal Canadian Cyber Incident Response Centre issued a security bulletin about the Heartbleed bug in April 2014 that affected many sites using specific versions of OpenSSL.

If vulnerabilities don't exist in the algorithm, attackers often look for weak keys or poorly implemented keys. For example, one of the core rules related to stream ciphers is to never reuse keys. If a protocol ignores this rule, the protocol is easier to crack. This emphasizes the importance of following fundamental key-management concepts.

Managing Cryptographic Keys

While SSCPs are not expected to design encryption algorithms, (ISC)² does expect you to understand some of the underlying concepts, especially related to the use of keys. This includes key rotation and key composition.

- **Key rotation** Key rotation refers to periodically replacing an old key with a new key. This is similar in concept to changing a password regularly. The length of time between using a different key varies for different algorithms. However, the important point here is that keys should be changed regularly.
- **Key composition** Key composition refers to creating strong keys. Strong keys are random and of sufficient length that they cannot be guessed. The length of the key varies for different algorithms. For example, AES-256 uses 256-bit keys and is very strong. However, the use of 256-bit keys for RSA is very weak. Instead, strong RSA keys are typically 1,024 or 2,048 bits in length. Algorithms include code to ensure keys are random.

Known-Plaintext Attack

A known-plaintext (or known-text) attack is possible when the attacker has samples of both plaintext and ciphertext data. For example, if an attacker has plaintext from an encrypted message, the attacker can then use different methods to try to decrypt the ciphertext to the known text. If the known text is decrypted, the same method can be used to decrypt similar data.

The chosen-plaintext attack is a variant of the known-plaintext attack. In this attack, the cryptanalyst doesn't have access to all of the plaintext but either has access to some plaintext or is able to predict plaintext that is included in the ciphertext. For example, if messages are frequently preceded with the phrase "From:," then the analyst can assume that "From:" is the known plaintext and use this information in the attack.

Ciphertext-Only Attack

In contrast, a ciphertext-only attack is one where the attacker has only ciphertext for analysis without any useful information about the plaintext data. If the attacker is able to discover any of the plaintext by analyzing the ciphertext, the attack is successful.

This type of attack is usually successful only when weak cryptographic methods are used. For example, if the same key is used for too long a period, or the key isn't protected, or a weak key is used, it makes it much easier for an attacker to break the code. WEP was cracked using a ciphertext-only attack due to poor key-management techniques used with WEP (among other vulnerabilities).

Chapter Review

Cryptography provides several important security benefits. You can encrypt data to prevent the loss of confidentiality, and you can use hashing algorithms to help ensure integrity. Many encryption and hashing protocols also provide authenticity, providing assurances that data is coming from a known source and is valid or reliable.

Organizations use cryptographic methods to secure sensitive data. In general, all non-public data should be encrypted to help ensure its protection. Additionally, regulations require organizations to protect certain data, such as personally identifiable information (PII) and protected health information (PHI). Encryption is a primary method of doing so.

Hashing provides one-way encryption to create a fixed-length hash (a number) from a file or message of any size. The hash created will always be the same as long as the file or message has not been modified. By comparing hashes at different times (such as before and after downloading a file), you can tell whether the original data has been modified or has lost integrity. If the hashes are the same, the data has not lost integrity. Message Digest 5 (MD5) and Secure Hashing Algorithm 1 (SHA-1) have been popular hashing algorithms, but both have been cracked. NIST recommends the use of SHA-3, which can create hashes of 224, 256, 384, or 512 bits.

Encryption converts plaintext data into ciphertext data. The two primary categories of encryption are symmetric encryption and asymmetric encryption.

Symmetric encryption uses a single key to encrypt and decrypt data. Keys used with symmetric encryption need to be protected and changed often. If the key is not protected, an attacker can use it to decrypt the data.

Advanced Encryption Standard (AES) is a strong, efficient symmetric encryption algorithm. NIST identified it in 2002 as a standard to be used by the U.S. government agencies, and it has been adopted in many different commercial applications. AES can use key sizes of 128 bits (AES or AES-128), 192 bits (AES-192), or 256 bits (AES-256). AES is the designated replacement for Data Encryption Standard (DES). DES has been cracked and is not recommended for use today. Other symmetric encryption algorithms include 3DES, Blowfish, International Data Encryption Algorithm (IDEA), and RC4.

Asymmetric encryption uses two keys (a public key and a private key) created as matched pairs. Information encrypted by one of the keys can only be decrypted with the matching key. The public key is embedded in a certificate and freely shared, but the private key always remains private. RSA is used for asymmetric encryption and employs large prime numbers to create matching public and private keys. Transport Layer Security (TLS) uses asymmetric encryption to share a symmetric key privately, and then it uses symmetric encryption to encrypt session data.

Secure/Multipurpose Internet Mail Extensions (S/MIME) is a popular standard used to protect e-mail. It supports encryption and digital signatures. Digital signatures are created by hashing an e-mail message and then encrypting the hash with the sender's private key. The recipient uses the sender's public key to decrypt the hash, and if the decryption succeeds, the digital signature provides authentication, integrity, and nonrepudiation.

Two other encryption schemes include steganography and Internet Protocol security (IPsec). Steganography is the practice of hiding data within data. IPsec is used to secure data going over a network and uses Authentication Header (AH) and Encapsulating Security Protocol (ESP). AH provides authentication and integrity, and ESP provides authentication, integrity, and confidentiality.

A public key infrastructure (PKI) includes all the components needed to issue and manage certificates, which follow the X.509 standard. The certificate authority (CA) issues the certificates. Certificates have expiration dates, but the CA can revoke the certificate earlier if necessary, such as if the matching private key is compromised. The CA publishes a list of revoked certificates by serial number in a certificate revocation list (CRL). Clients can request a copy of the CRL to validate a certificate. The Online Certificate Status Protocol (OCSP) is an alternative method used to validate certificates. The first CA in a certificate trust chain is called the root CA. If the root CA is trusted, all certificates issued by any CAs in the trust path are also trusted.

Cryptanalysis is the process of deciphering codes through analysis. Two common cryptanalysis attacks are a known-plaintext attack and a ciphertext-only attack. A known-plaintext (or known-text) attack is possible when the attacker has samples of both plaintext and ciphertext data. A ciphertext-only attack is one where the attacker has only ciphertext for analysis.

Questions

1. What basic security function does asymmetric encryption provide?
 - A. Integrity
 - B. Authentication
 - C. Confidentiality
 - D. Availability
2. Which of the following choices allows you to verify that a file has not been modified?
 - A. AES
 - B. SHA
 - C. PKI
 - D. IDEA
3. Which of the following choices provides one-way encryption of data?
 - A. Symmetric
 - B. Asymmetric
 - C. Transport Layer Security
 - D. Hashing
4. Which of the following uses a single key to encrypt and decrypt data?
 - A. Symmetric
 - B. Asymmetric
 - C. Public key cryptography
 - D. SHA-1
5. Which of the following keys is changed the most often?
 - A. Public key
 - B. Private key
 - C. Symmetric key
 - D. Session key
6. Which of the following is NOT a symmetric encryption standard?
 - A. AES
 - B. Blowfish
 - C. RC4
 - D. RSA

7. Which of the following is a symmetric 128-bit block cipher?
 - A. Data Encryption Standard (DES)
 - B. Triple Data Encryption Standard (3DES)
 - C. Advanced Encryption Standard (AES)
 - D. Blowfish
8. Of the following choices, which one is considered a strong, efficient symmetric encryption algorithm?
 - A. TLS
 - B. DES
 - C. 3DES
 - D. AES
9. What type of cryptography does public cryptography use?
 - A. Asymmetric encryption
 - B. Symmetric encryption
 - C. Steganography
 - D. One-way functions
10. Which of the following is an accurate statement related to asymmetric encryption?
 - A. It is used to privately share a private key.
 - B. It is used to privately share a public key.
 - C. It is used to privately share a secret key.
 - D. It is faster than symmetric encryption.
11. What is a common standard used to encrypt and digitally sign e-mail?
 - A. Symmetric encryption
 - B. S/MIME
 - C. TLS
 - D. Steganography
12. Of the following choices, what is NOT provided with a digital signature used for e-mail?
 - A. Authentication
 - B. Integrity
 - C. Confidentiality
 - D. Nonrepudiation

13. What is used to create a digital signature used with e-mail?
 - A. The public key of the sender
 - B. The private key of the sender
 - C. The public key of the recipient
 - D. The private key of the recipient
14. Which of the following would most likely be used to encrypt data in an e-mail message before it is sent?
 - A. The public key of the sender
 - B. The private key of the sender
 - C. The public key of the recipient
 - D. The private key of the recipient
15. Someone has embedded a secret code within a picture used on a web page. What is the best description of this?
 - A. Symmetric encryption
 - B. Asymmetric encryption
 - C. Hashing
 - D. Steganography
16. How are public keys distributed to clients from Internet websites?
 - A. As e-mail attachments
 - B. Embedded in certificates
 - C. As cookies
 - D. Embedded in the HTML code for the page
17. How are public keys shared with other entities?
 - A. Published in a certificate
 - B. Encrypted by a private key
 - C. Encrypted by a session key
 - D. Public keys are not shared
18. A website sent a user a certificate to initiate a secure web session over the Internet. What information would NOT be in the certificate?
 - A. Name of the website
 - B. Name of the issuing CA
 - C. Private key
 - D. Expiration date

19. Of the following choices, what is used to determine whether a certificate has been revoked?
 - A. OCSP
 - B. Digital signature
 - C. CARL
 - D. Trust chain
20. Researchers are attempting to discover weaknesses in an encryption algorithm using a known-plaintext attack. What is this called?
 - A. Cryptography
 - B. Cryptanalysis
 - C. Criminal behavior
 - D. Hashing

Answers

1. **C.** Encryption (any type of encryption, including both asymmetric and symmetric encryption) provides confidentiality for data. Hashing methods provide integrity. Authentication proves the identity of a user or system. Availability ensures that systems and data are available when needed.
2. **B.** Secure Hashing Algorithm (SHA) is a hashing algorithm, and hashing is a key method of ensuring integrity (or verifying a file has not been modified). The hash is calculated at two different times, and if the hash is the same, the file has not been modified. Advanced Encryption Standard (AES) is a strong symmetric encryption protocol. A public key infrastructure (PKI) is used to support the creation, management, and distribution of certificates. International Data Encryption Algorithm (IDEA) is an older symmetric encryption protocol.
3. **D.** Hashing algorithms use a one-way encryption method to create a hash from the data. They are also known as one-way functions. The hashing algorithm creates a fixed-length hash from a file, but the hash cannot be used to re-create the file. Symmetric encryption and asymmetric encryption are both two-way, because data can be encrypted and decrypted. Transport Layer Security (TLS) is widely used to encrypt Internet traffic and uses both symmetric and asymmetric encryption.
4. **A.** Symmetric encryption uses a single key to encrypt and decrypt data. Asymmetric encryption uses two keys (a public key and a private key) to encrypt and decrypt information and is often referred to as public key cryptography. Secure Hashing Algorithm 1 (SHA-1) is a hashing algorithm and it doesn't use a key.
5. **D.** A session key is only used for a session (such as a web browsing session) and is changed more often than the other keys. Public and private keys typically last for a year or longer. Symmetric encryption uses a symmetric key (also called a secret key), which can stay the same for a specific piece of data as long as the data remains encrypted.

6. **D.** RSA is an asymmetric encryption standard using public and private keys and is widely used with Transport Layer Security (TLS). The other choices are all symmetric encryption standards using a single key to encrypt and decrypt the data.
7. **C.** AES is a 128-bit block cipher. All of the other answers are 64-bit block ciphers.
8. **D.** The Advanced Encryption Standard (AES) is considered a strong, efficient symmetric encryption algorithm and it is widely used. DES is an older algorithm that has been cracked. 3DES is strong, but takes more processing power and is less efficient than AES. Transport Layer Security (TLS) uses both symmetric and asymmetric encryption, and calling it a symmetric encryption algorithm is inaccurate.
9. **A.** Public key cryptography uses asymmetric encryption with two matched keys (a public key and a private key) to encrypt and decrypt information. Symmetric encryption uses a single key (often called a session key) to encrypt and decrypt data. Steganography hides data within data. Hashes are also known as one-way functions and they provide integrity.
10. **C.** Asymmetric encryption is used to privately share a secret key (or session key). Asymmetric encryption uses a matched pair of keys known as a private key and a public key. The private key is never shared, and the public key is publicly shared in a certificate. Symmetric encryption is faster than asymmetric encryption.
11. **B.** Secure/Multipurpose Internet Mail Extensions (S/MIME) is the standard used to encrypt and digitally sign e-mail. Symmetric encryption uses a single key to encrypt and decrypt data, but cannot digitally sign e-mail. Transport Layer Security (TLS) encrypts data sent over a network and is used with HTTPS. Steganography is the practice of hiding data within data or in plain sight.
12. **C.** A digital signature does not provide confidentiality because the digital signature does not encrypt the data. A digital signature does provide authentication, integrity, and nonrepudiation. It's possible to digitally sign an e-mail without encrypting it.
13. **B.** A digital signature is created by hashing a message and encrypting the hash with the sender's private key. The recipient can then decrypt the hash with the sender's public key. The recipient's keys are not used for a digital signature, but they are used to encrypt and decrypt e-mail.
14. **C.** E-mail is encrypted using the recipient's public key. The recipient's public key actually encrypts a symmetric key and uses the symmetric key to encrypt the e-mail. The recipient uses the recipient's private key to decrypt the symmetric key and then decrypts the message with the symmetric key. The sender's keys are not used to encrypt or decrypt e-mail.
15. **D.** Steganography is the practice of hiding data within data, such as embedding a secret code within a picture. Symmetric encryption uses a single key for encryption and decryption of data, while asymmetric encryption uses two keys (a public key and a private key) for encryption and decryption. Hashing creates a hash that can be used for integrity.

16. **B.** Public keys are embedded in certificates and distributed to clients in the certificate. Although users can send certificates to each other as e-mail attachments, a website does not use this method. Public keys are not included in cookies or in HTML code.
17. **A.** Public keys are published in certificates. They are never encrypted. They are shared so that other entities can use them for asymmetric encryption.
18. **C.** The private key is not included in the certificate but instead is kept private on the server. The public key is included in the certificate along with the name of the website, the name of the CA that issued the certificate, the expiration date of the certificate, and more.
19. **A.** The Online Certificate Status Protocol (OCSP) is used to verify the health of a certificate. An OCSP responder will indicate whether a certificate has been revoked when queried with the certificate's serial number. A digital signature uses certificates but doesn't determine whether a certificate is revoked. CAs issue a certificate revocation list (CRL), but CARL isn't a valid acronym in the context of checking certificates. A trust chain determines if the CA that issued the certificate is trusted, but doesn't indicate if a certificate is revoked.
20. **B.** Cryptanalysis is the process of deciphering codes through analysis, and a known-plaintext attack is one method of cryptanalysis. Cryptography is the science of using different methods and techniques to encrypt data. It is not criminal to search for weaknesses, but the action taken after these weaknesses are discovered can be criminal. Hashing is the process of creating a hash from a file or a message with a hashing algorithm, and it is used to prevent the loss of integrity.

This page intentionally left blank

About the Download

This e-book comes complete with Total Tester customizable practice exam software.

System Requirements

The Total Tester 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.

Downloading Total Tester Premium Practice Exam Software

To download the Total Tester software, simply click the link below and follow the directions for free online registration.

www.totalsem.com/1259583066dl

Total Tester Premium Practice Exam Software

Total Tester provides you with a simulation of the actual exam. You can also create custom exams from selected certification domains 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 domain.

To take a test, launch Total Tester 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

Once you've downloaded the Total Tester software, double-click the Launch.exe icon. From the main screen you may install the Total Tester software by clicking the Total Tester Practice Exams link. 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.

Technical Support

Technical Support information is provided as follows, by feature.

Total Seminars Technical Support

For questions regarding the Total Tester software, visit www.totalsem.com or e-mail support@totalsem.com.

McGraw-Hill Education Content Support

For questions regarding e-book content, e-mail customer.service@mheducation.com. For customers outside the United States, e-mail international_cs@mheducation.com.

3DES A symmetric encryption standard. It improves Data Encryption Standard (DES) by encrypting data in three passes with three separate keys. It was one of the standards evaluated by the National Institute of Standards and Technology (NIST) with Advanced Encryption Standard (AES), but was not selected. It is a slower and processor-intensive block cipher, but is still strong and used in some applications. It is also called *triple DES* and *three DES*.

3G Third generation of wireless technologies. 1G uses analog signals, and 2G uses digital signals. 3G provides higher transfer speeds using digital technologies.

4G Fourth generation of wireless technologies. It's an improvement over 3G and includes speeds as high as 100 Mbps for highly mobile users in trains and cars, and 1 Gbps for individuals walking or staying still.

AAAs of security Authentication, authorization, and accounting. Authentication ensures that entities can prove their identity. Authorization grants access to individuals based on their proven identity. Accounting tracks and records their activity in logs.

access control Mechanism used to restrict or control access to resources. Access controls can be logical (implemented with technology) and implemented by the security kernel, or physical, such as a locked door or security guard. Access controls allow subjects access to objects, such as allowing a user to access a file. Some relevant access control models are Mandatory Access Control (MAC), Discretionary Access Control (DAC), Role-based Access Control (Role-BAC), Rule-based Access Control (Rule-BAC), and Attribute-based Access Control (ABAC). Biba, Bell-LaPadula, Clark-Wilson, and the Chinese Wall are specific MAC models.

access control list (ACL) A list of rules. ACLs are most commonly associated with routers and firewalls. The rules identify the traffic allowed in or out of a network. On host-based firewalls, the rules identify the traffic allowed in or out of a system.

accountability The ability of a system to track the activity of an individual. It depends on proper identification and authentication. If a system can identify individual users, track their actions, and monitor their behavior, it provides accountability.

accounting Logs that track the activity of a user through monitoring. One method of accounting is audit logs that create an audit trail.

accreditation The process of formally declaring that the system is approved to operate. Accreditation comes after system certification.

Address Resolution Protocol (ARP) Protocol that resolves Internet Protocol (IP) addresses to Media Access Control (MAC) addresses. ARP is used on the Data Link layer of the Open Systems Interconnection (OSI) Model. Compare to *Bootstrap Protocol (BootP)* and *Reverse Address Resolution Protocol (RARP)*.

Advanced Encryption Standard (AES) A strong, efficient symmetric encryption algorithm. The National Institute of Standards and Technology (NIST) selected it in 2002 as a replacement for Data Encryption Standard (DES) as the standard used by the U.S. government. AES has since been adopted in both the public and private sectors and is widely used today in many applications. AES uses key sizes of 129, 192, and 256 bits and sometimes is listed based on the key size (such as AES-256).

advanced persistent threat (APT) A highly sophisticated group of attackers who have the capability and intent to carry out successful attacks. Many governments are suspected of sponsoring groups known as APTs.

anomaly based A method of detection used by intrusion detection systems (IDSs) and intrusion prevention systems (IPSs). The IDS/IPS attempts to document normal behavior in the form of a baseline. It then monitors the activity and constantly compares it to the baseline. If the current activity differs significantly from the baseline, the IDS/IPS will issue an alert on the activity.

antivirus (AV) software A primary method used to detect and prevent infections from malware. In addition to being able to detect and prevent infections, most AV software is able to remove the malware, restore the infected file to its original state, or quarantine the file.

armored virus A virus that uses code to make it difficult for AV researchers to reverse-engineer the code. Encryption is often combined with other methods to prevent reverse-engineering.

asymmetric encryption A process of encrypting and decrypting data using two matched keys known as a public key and a private key. It is also known as *public key cryptography*. Anything encrypted with the public key can be decrypted only with the matching private key. Anything encrypted with the private key can be decrypted only with the matching public key. The private key is always kept private and never shared. The public key is freely shared and publicly available.

attribute In a database, a table column. Data within the table is stored in rows, or tuples.

Attribute-based Access Control (ABAC) An access control model that uses attributes to determine access. It evaluates subject and object attributes and grants access based on the value of these attributes. Attributes can be almost any characteristic of

a user, the environment, or the resource. Other access control models are Mandatory Access Control (MAC), Discretionary Access Control (DAC), Role-based Access Control (Role-BAC), and Rule-based Access Control (Rule-BAC).

audit trail A record of events occurring on a system or network, recorded in one or more logs. When you have access to all the logs, you are able to re-create the events that occurred leading up to an event and identify what actually occurred during an event.

authentication Proof of an identity, established by providing credentials. There are three types or factors of authentication: something you know (such as a username and password), something you have (such as a smart card), and something you are (using biometrics).

Authentication Header (AH) Used with Internet Protocol security (IPsec) to provide authentication between systems and integrity of the packets. The AH is created from a hash of the packet with additional authentication data, and the hash is encrypted to prevent tampering. The hash provides integrity. Packets protected with AH are identified with protocol ID 51. See also *Encapsulating Security Protocol (ESP)* and *Internet Protocol security (IPsec)*.

authenticity Assurance that data is coming from a known source and is valid or reliable. Validation processes, such as many encryption and hashing protocols, provide authenticity. Additionally, some processes, such as digital signatures, have authenticity as a key goal.

authorization A process that provides access to resources through the assignment of permissions. This process starts with authentication, where users claim an identity and prove their identity with a password or other credentials. Once users have been authenticated, authorization defines the resources that a user can access and the rights that a user can invoke.

availability One of the three main goals of information security known as the CIA security triad. Availability ensures that systems and data are up and operational when needed. The other two parts of the CIA triad are confidentiality and integrity.

backdoor A nontraditional method of accessing an application or system. It can be code embedded in an application that provides access to the application, the application's code, or data via a covert method. Attackers often try to install a backdoor onto a system after infecting it with malware that grants them remote access.

baseline A known starting point. Baselines are an important element of configuration control and are often implemented with images. If the baseline configuration is known, it's relatively simple to check the system to determine whether the configuration has been modified from the baseline. Anomaly-based intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) also use baselines by first documenting normal behavior in the form of a baseline. The system then monitors the activity and constantly compares it to the baseline.

bcrypt A key-stretching algorithm used to protect passwords on UNIX and Linux distributions stored in the shadow password file. It salts passwords before encrypting them with Blowfish, which helps thwart rainbow table attacks.

Bell-LaPadula model An access control model used to ensure confidentiality. It uses two primary rules: no read up and no write down. Compare to *Biba model*.

Berkeley Internet Name Domain (BIND) A version of DNS software that runs on UNIX systems. It is freely available and runs on many DNS servers on the Internet.

Biba model An access control model used to ensure integrity. It uses two primary rules: no read down and no write up. Compare to *Bell-LaPadula model*.

biometrics Method of identifying unique characteristics of a person, such as a fingerprint or retina scan. Biometrics provide authentication in the “something you are” factor.

bit-copy tools Software used to capture the contents of a drive without modifying the data. The copy is known as a forensic duplicate image and can be examined without affecting the original.

blacklist A list of prohibited applications. Application blacklisting identifies specific applications that cannot run on a system. Compare to *whitelist*.

Bootstrap Protocol (BootP) A protocol that provides an IP address to clients and can be used to retrieve a bootable image for clients. It is similar to Reverse Address Resolution Protocol (RARP), although RARP only retrieves the IP address.

botnet A group of computers (called zombies) controlled by an attacker. The term *botnet* is derived from *robot* and *network*. The attacker manages a command control center, and the computers in the botnet do the bidding of the attacker.

bring your own device (BYOD) An organizational policy that allows users to bring their personally owned devices (such as smartphones and tablets) to work and connect them to the organization’s network. Compare to *corporate-owned, personally enabled (COPE)* policies.

buffer overflow An error that can occur when a system receives more data than it expects and is unable to handle it gracefully. Attackers attempt to exploit buffer overflow errors to install malware on systems.

buffer overflow attack An attack on a system that has a buffer overflow vulnerability. Buffer overflow vulnerabilities can be reduced with input validation techniques and by keeping systems up to date.

bus Network configuration where all computing devices are connected directly to each other via a shared cable connection. Both ends of the bus must be terminated. If one of the terminators is not present or the cable is disconnected, communication with all devices on the bus stops. See also *mesh, star, token ring, and tree*.

business continuity plan (BCP) A written document that includes the processes and procedures to prevent mission-critical services from being interrupted or disrupted. A BCP includes disaster recovery elements used to restore the organization to fully functioning operations as quickly and efficiently as possible.

business impact analysis (BIA) A part of a BCP. It identifies the impact to the organization if any business functions are lost due to any type of incident. It helps an organization identify what business functions are critical to continued operations by identifying the impact to the business if a business function stops.

certificate A file used for security purposes, such as authentication, encryption, protection of e-mail, and code signing. Certificate authorities issue and manage certificates.

certificate authority (CA) An entity that issues and manages certificates through their lifetimes. CAs can be public (such as VeriSign) or private within an organization. Public CAs sell and validate certificates and provide assurances to users that certificates are valid. A CA is sometimes referred to as a *certification authority*.

certification A process that evaluates, describes, and tests a system and all of the controls that are in place to mitigate risks to the system. After a system has been certified, a separate accreditation process formally approves the system to operate.

chain of custody A document that shows exactly where a piece of evidence is from the point it is collected until it is disposed. A chain of custody provides proof that evidence has been protected. If the chain-of-custody document is not present, the validity of the evidence can be questioned and its usefulness negated.

change control A process that helps prevent unintended outages from changes. A change control process gives experts an opportunity to examine the change for potential problems before the change is implemented. It is also called *change management*.

Chinese Wall model An access control model used to help prevent a conflict of interest. Data is classified based on conflict-of-interest classes. Users who have access to one class are denied access to data in conflicting classes. Also known as the *Brewer-Nash model*.

CIA security triad The three main goals of information security: confidentiality, integrity, and availability. Confidentiality ensures that unauthorized individuals cannot access data. Integrity prevents any unauthorized or unwanted modification of data or system configurations. Availability ensures that systems and data are available when needed.

ciphertext Data that is encrypted and not readable until it is decrypted into plaintext. Authorized personnel can decrypt ciphertext to return it to plaintext and read it. Ideally, unauthorized personnel cannot read ciphertext. Encryption techniques convert plaintext to ciphertext.

ciphertext attack A cryptanalysis attack. Attackers use this when they only have ciphertext for analysis without any useful information about the plaintext data.

Clark-Wilson model An access control model used to ensure integrity; the model has stricter rules than the Biba model. These integrity rules enforce the principle of separation of duties.

cleartext Data that is not in an encrypted format. It is easily readable without any cryptographic or cryptanalysis techniques. Encryption techniques convert plain text to ciphertext. Also called *plain text*.

clipping level A predetermined threshold level. Many automated accounting systems use clipping levels to generate alerts after the system detects a preset number of events. An auditing system ignores events until the number of events reaches the clipping level.

code of ethics Statements and principles that individuals can use to guide their decisions and help in ethical dilemma situations. (ISC)²'s Code of Ethics includes a preamble and four canons describing ethical expectations from its certified practitioners. Candidates must commit to and abide by them to earn and keep the SSCP certification. Organizations can also use ethics statements for internal employees.

cold site An alternative location used in business continuity planning. A cold site is a building with a roof, running water, and electricity. It doesn't include the necessary hardware, software, or personnel. In the event of an emergency, personnel move all of the resources to the cold site location, hook them up, and configure them for operation. Compare to *hot site*, *warm site*, and *mobile site*.

command injection An attack that attempts to inject commands into an application. In some cases, a command injection attack can inject operating system commands that would normally be executed at the command line. In other cases, it injects code such as JavaScript or SQL statements. Input validation techniques help mitigate command injection attacks.

Common Criteria A framework used to evaluate systems, formally known as Common Criteria for Information Technology Security Evaluation. It provides assurances that the specification, implementation, and evaluation of a system's security has gone through a rigorous and standardized process.

community cloud A cloud-based service shared by two or more organizations. It is similar to a private cloud in that it is not available to the public. Compare to *public cloud*, *private cloud*, and *hybrid cloud*.

confidentiality Ensures that unauthorized entities cannot access data. Access controls and encryption help protect against the loss of confidentiality. Confidentiality is one of the three main goals of information security known as the CIA security triad. The other two goals are integrity and availability.

configuration management A process that ensures that information about system configuration is available for any system and helps ensure that similar systems are configured similarly.

control A means, method, action, technique, process, procedure, or device that reduces the vulnerability of a system or the possibility of a threat exploiting a vulnerability in a system. Controls are risk management tools. The terms *control*, *countermeasure*, and *safeguard* are often used interchangeably.

corporate-owned, personally enabled (COPE) An alternative to bring your own device (BYOD) policies. An organization purchases and issues devices (such as smartphones and tablets) to users instead of allowing them to connect their personally owned devices to the network. Compare to *bring your own device (BYOD)* policies.

countermeasure A means, method, action, technique, process, procedure, or device that reduces the vulnerability of a system or the possibility of a threat exploiting a vulnerability in a system. Controls are risk management tools. The terms *control*, *countermeasure*, and *safeguard* are often used interchangeably.

Crossover Error Rate (CER) The point where the False Acceptance Rate (FAR) and the False Rejection Rate (FRR) of a biometric system are equal or cross over. A lower CER indicates a better-performing biometric system. Compare to *False Acceptance Rate (FAR)* and *False Rejection Rate (FRR)*.

cross-site request forgery (CSRF or XSRF) An attack that results in websites executing unauthorized commands as if a user requested them. Attackers create specially crafted links and encourage users to click the links. For example, attackers can send the links in phishing e-mails.

cross-site scripting (XSS) An attack that attempts to inject HTML or JavaScript into a web page. After the attack, the code executes on a user's system when the user visits the attacked website. Successful attacks sometimes allow attackers to read cookies and use this data to launch session-hijacking attacks.

cryptanalysis The process of deciphering codes through analysis. Both attackers and researchers use cryptanalysis techniques. Researchers use them to search for weaknesses in cryptography, with the goal of improving cryptographic methods. Attackers use them to search for weaknesses, with the goal of exploiting them.

cryptography The science of scrambling, or encrypting, data. Cryptography includes the study of algorithms and other methods that help protect the confidentiality of data.

data at rest Any data that is in computer storage, such as on system hard drives, portable USB drives, flash drives, storage area networks, and backup tapes. Compare to *data in motion*.

Data Encryption Standard (DES) An older symmetric encryption standard using 56-bit keys. It has been broken and is rarely used anymore.

data in motion Any data being transmitted over a network (sometimes called *data in transit*). It includes data transmitted over an internal network using wired or wireless methods, as well as data transmitted over the Internet. Compare to *data at rest*.

data inference The ability of someone to gain knowledge by piecing together unclassified data to determine classified or secret information. It also includes the ability to gather large quantities of information and attempt to learn details of the information through deduction.

data loss prevention (DLP) Techniques used to monitor data usage and prevent the unauthorized use or transmission of sensitive data. Different types of DLP systems can monitor data in motion and data at rest.

database view A virtual table that provides access to specific columns in one or more tables. A view doesn't hold any data but presents the data in the underlying table or tables. A database administrator can grant access to a view without granting access to a table to limit what a user can see and manipulate.

decryption The process of converting ciphertext data into plaintext data. Data is encrypted to prevent loss of confidentiality. Compare to *encryption*.

deduplication The process of keeping only a single copy of a file on a system instead of multiple identical files. Deduplication saves storage space.

defense diversity A defense-in-depth strategy using dissimilar technologies. Implementing a demilitarized zone (DMZ) with firewalls from two separate vendors is an example of defense diversity.

defense in depth A strategy that provides a layered approach to security. Instead of using one or two security controls, multiple controls are used. If one control fails, other controls continue to provide protection.

demilitarized zone (DMZ) A perimeter network used to host resources on the Internet (such as web servers, e-mail servers, or FTP servers). The DMZ provides a layer of protection for the resources that would not be available if they were placed directly on the Internet.

denial of service (DoS) An attack that attempts to prevent a system from answering legitimate requests from users, directly affecting the availability portion of the CIA triad. The attack is launched by a single system.

differential backup A type of backup used in full/differential backup strategies. Differential backups only back up data that has changed since the last full backup, without regard to any other differential backups. Compare to *incremental backup*.

digital signature A file used to provide authentication, integrity, and nonrepudiation security for e-mail. A digital signature is created by hashing an e-mail message and then encrypting the hash with the sender's private key. It is decrypted with the sender's public key.

disaster recovery plan (DRP) A document used to provide an organization with a plan to restore critical operations after a disaster. The overall goal is to provide employees with clear-cut steps on what to do and the order of these steps.

Discretionary Access Control (DAC) model An access control model commonly used to control access for file systems, such as New Technology File System (NTFS) or Network File System (NFS). Objects (such as files and folders) are owned by users. Users have full control over the objects and can grant others access. DAC provides the most granular level of control. Other access control models are Mandatory Access Control (MAC), Role-based Access Control (Role-BAC), Rule-based Access Control (Rule-BAC), and Attribute-based Access Control (ABAC).

distributed denial of service (DDoS) A DoS attack launched against a single system from multiple attackers. Botnets are often used in DDoS attacks.

Domain Name System (DNS) Provides name resolution services by resolving host names to IP addresses and IP addresses to host names. DNS uses port 53.

dropper A type of Trojan horse that attempts to drop malware onto a computer. It might be from an infected file delivered via e-mail or delivered via a drive-by download.

due care The practice of implementing security policies and practices to protect resources. Organizations are required by law to exercise due care or they can be found legally negligent.

due diligence The investigative steps that an organization takes prior to taking on something new, such as signing a contract or making a major purchase. An organization has an obligation to exercise due diligence to discover risks.

Dynamic Host Configuration Protocol (DHCP) A protocol used to assign TCP/IP configuration information to DHCP clients. This includes an IP address, a subnet mask, the address of a DNS server, and much more. DHCP uses UDP ports 67 and 68.

Encapsulating Security Protocol (ESP) One of the underlying protocols of IPsec, ESP encrypts the data and also provides the same authentication services provided from Authentication Header (AH). Data encrypted with ESP cannot be read if captured with a sniffer. Packets protected with ESP are identified with protocol ID 50.

encryption The process of converting plaintext data into ciphertext data to prevent loss of confidentiality. The process is reversed by decrypting the ciphertext data to create the original plaintext data. Compare to *decryption*.

Ethernet Set of local area network (LAN) standards defined by the IEEE 802.3 documents. It is the most widely used standard for LANs.

extranet An area of an organization's network used to host resources via the Internet but is available only to trusted entities. An extranet is available via the Internet, but only to a specific target audience.

False Acceptance Rate (FAR) A type 2 biometric error. It refers to the percentage of times a biometric system falsely identifies an unknown user as a known user. Compare to *False Rejection Rate (FRR)* and *Crossover Error Rate (CER)*.

false positive A false indication of an attack or a vulnerability. Vulnerability scanners can give false positives indicating that a vulnerability exists even though it doesn't. Similarly, intrusion detection systems can give an indication that a system is being attacked even when it isn't.

False Rejection Rate (FRR) A type 1 biometric error. It refers to the percentage of times a biometric system falsely rejects a known user and instead indicates that the user is unknown. Compare to *False Acceptance Rate (FAR)* and *Crossover Error Rate (CER)*.

fault tolerance The practice of adding and maintaining redundancy for disks, servers, connections, and sites. Fault tolerance controls help eliminate outages from single points of failure by adding redundancies.

federated access A single sign-on (SSO) technology that allows users in different networks to access multiple systems after logging on once. The systems can be using different operating systems owned and managed by different organizations.

Fiber Distributed Data Interface (FDDI) An extension of token ring networks that uses fiber-optic connections instead of copper. FDDI uses dual rings, and the second ring provides redundancy. FDDI token rings support speeds up to 100 Mbps over distances as far away as 120 miles.

file integrity checker An application that can verify that files have not been modified. In doing so, the application guards against a loss of integrity. File integrity checkers use hashing algorithms to capture hashes of files in a known good state. Later, they create hashes on the same file and compare the two hashes. If the file is unchanged, the hashes are the same, but if the file has been modified, the hashes are different.

File Transfer Protocol (FTP) A protocol used to transfer files from one system to another. Data is transmitted in cleartext using TCP ports 20 and 21. See also *Secure File Transfer Protocol (SFTP)* and *Trivial File Transfer Protocol (TFTP)*.

firewall Software or hardware used to filter traffic into or out of a network. A firewall can be a dedicated physical device or an additional application running on a system such as a desktop computer.

first responder The first forensics-trained person on the scene of an information technology (IT) incident. It could be a highly trained forensics expert or a system administrator with only basic training. One of the primary responsibilities of a first responder is the preservation of the scene of the incident and all associated evidence.

foreign key A type of key used in databases. A foreign key in one table points to a unique primary key in another table to create a relationship between the two tables.

forensics The science of examining and inspecting crime scenes for evidence. Computer forensics is the science of examining and inspecting computer systems for evidence about an event or crime.

guidelines Recommendations provided to members of an organization. Guidelines aren't mandatory or authoritative in nature. Guidelines (and procedures) are derived from policies.

hardening The practice of making a system more secure from its default configuration. Hardening often includes removing or disabling unused protocols and services, changing defaults, keeping systems up to date, enabling firewalls, and using AV software.

hardware token A device held by a user that displays a number or a password that changes frequently, such as every 60 seconds. The number is synchronized with a server and used as a one-time password.

hashing

The process of creating a hash (a number) by executing a hashing algorithm against a piece of data. As long as the original data is not changed, the hash will always be the same. You can execute the hashing algorithm on the file or message at one point to create a hash and then later execute the same hashing algorithm again. If the hashes are the same, the original data has not been modified and it has retained integrity.

heuristics-based detection A method of virus detection used to detect previously unknown viruses. Heuristics attempts to detect a virus based on its behavior.

honeypot A computer set up to entice would-be attackers. It is often configured with weak security so that an attacker can easily hack into it, and usually has fake data that has no use to the company.

host-based firewall A firewall that provides protection for a single host. Many operating systems include host-based firewalls running as additional software. Compare to *network-based firewall*.

host-based IDS (HIDS) An intrusion detection system (IDS) installed on an individual system, such as a server or workstation. It can only monitor activity on the host. Compare to a *network-based IDS (NIDS)*, which monitors overall network activity.

hot site An alternative location used in business continuity planning. A hot site includes all of the necessary resources to take over the operations of another location in a very short period of time, sometimes within minutes. It includes hardware such as servers and the network infrastructure, up-to-date data, and personnel to manage the functions of the alternative location. Compare to *cold site*, *warm site*, and *mobile site*.

hybrid cloud A cloud-based service that is a combination of any two or more clouds. Compare to *public cloud*, *community cloud*, and *private cloud*.

HyperText Transfer Protocol (HTTP) The common protocol used to transfer web pages over the Internet. Most web browsers use HTTP by default. HTTP uses TCP port 80.

HyperText Transfer Protocol Secure (HTTPS) An encrypted form of HTTP. Transport Layer Security (TLS) replaced Secure Sockets Layer (SSL) in most implementations of HTTPS. HTTPS uses TCP port 443.

identification The act of claiming an identity, typically by using a username. Users prove identity by providing proper credentials such as a password.

identity management The process of managing accounts and access to resources. Provisioning includes creating accounts for users and granting appropriate privileges. Maintenance ensures that password policies and account policies are implemented properly and ensures that inactive accounts are disabled. Entitlement helps ensure that the principle of least privilege is enforced.

incident Any violation of policies or security practices that has the potential to result in an adverse event or that has resulted in an adverse event. Risk management practices attempt to prevent any incidents, detect incidents when they occur, and correct problems that occur from incidents. Also called a *security incident*.

incident response The process of responding to an incident. Many organizations have formal procedures for responding to incidents.

incremental backup A type of backup used in full/incremental backup strategies. Incremental backups back up only data that has changed since the last full or incremental backup. Compare to *differential backup*.

information rights management (IRM) Methods used to protect documents after they've been sent to an untrusted party. IRM can prevent recipients from forwarding, copying, modifying, printing, faxing, or pasting the contents from a file.

Infrastructure-as-a-Service (IaaS) A type of cloud computing that provides customers with hardware (such as servers) and infrastructure (such as routers or switches). Customers maintain all software on the hardware, but the provider owns and maintains the hardware. Also known as Hardware-as-a-Service. Compare to *Software-as-a-Service (SaaS)* and *Platform-as-a-Service (PaaS)*.

input validation The process of checking data before using it within an application. Web applications use input validation to prevent different types of injection attacks, such as SQL injection or cross-site scripting attacks. Input validation also helps prevent buffer overflow attacks.

integrity Assurance that data or a system configuration has not been modified. Hashing and audit logs are two methods used to ensure integrity. Integrity is one of the three main goals of information security known as the CIA security triad. The other two goals are confidentiality and availability.

Internet Control Message Protocol (ICMP) A protocol used to check or verify the health of a network or network device. ICMP is used with many diagnostic protocols such as ping, pathing, and tracer.

Internet Group Message Protocol (IGMP) A protocol used for IPv4 multicasting, where data is sent from one computer to multiple computers at the same time. IPv6 uses ICMPv6 messaging instead of IGMP.

Internet Message Access Protocol version 4 (IMAP4) A protocol used to manage e-mail in folders. IMAP4 servers store e-mail on the server and allow users to organize the e-mail in folders. Compare to *Post Office Protocol version 3 (POP3)* and *Simple Mail Transfer Protocol (SMTP)*.

Internet Protocol security (IPsec) A security protocol used to provide security for IP traffic traveling over a network. IPsec includes Authentication Header (AH) and Encapsulating Security Protocol (ESP). AH provides authentication between the systems and verifies the integrity of the packets. ESP encrypts the data and provides the same authentication services provided by AH.

intranet A network that is internal to an organization. An intranet has private IP addresses, and clients within the intranet access the Internet with Network Address Translation (NAT), which translates private IP addresses to public and public IP addresses to private.

intrusion detection system (IDS) A system designed to provide continuous monitoring of networks and hosts to help protect them from attacks. The goal is to detect an attack as it is occurring. Some IDSs are passive and will provide a notification of a potential attack, and other IDSs are active and will thwart the attack in progress.

intrusion prevention system (IPS) Similar to an intrusion detection system (IDS), but with two primary differences. An IPS is placed inline with the traffic, preventing the malicious traffic from reaching the protected network. An IPS provides an active response.

job rotation The practice of moving employees between different jobs on a periodic basis. Job rotation is a security practice that can reduce the risk of fraud within a company. Organizations commonly combine it with a separation of duties policy.

key stretching A method that transforms a potentially weak password or encryption key into a more secure password or encryption key. Password-Based Key Derivation Function 2 (PBKDF2) and bcrypt are two key-stretching algorithms.

known-plaintext attack A cryptanalysis attack used when the attacker has samples of both plaintext and ciphertext data. For example, if an attacker has plaintext from an encrypted message, the attacker can then use different methods to try to decrypt the ciphertext to the known text. If the known text is decrypted, the same method can decrypt similar data.

Layer 2 Tunneling Protocol (L2TP) A tunneling protocol used with some virtual private networks (VPNs). It combines the strengths of Layer 2 Forwarding (L2F) and Point-to-Point Tunneling Protocol (PPTP) and is documented in RFC 2661. L2TP VPN traffic is commonly encrypted with IPsec (as L2TP/IPsec).

least privilege An important security principle that ensures that users are given only the rights and privileges that they need to perform their job, and no more.

logic bomb Malicious code that executes in response to an event, such as a point in time, or in response to specific actions.

malware Malicious software or malicious code. It includes any code or software that can be described as being harmful or destructive to computers, networks, or the computing environment as a whole. Some of the common malicious code types include viruses, worms, logic bombs, and Trojan horses.

Mandatory Access Control (MAC) model An access control model that provides the highest level of security. It is used by the U.S. military. Subjects and objects are assigned labels; subjects (such as users) are able to access objects (such as files) only when the labels match. Other access control models are Discretionary Access Control (DAC), Role-based Access Control (Role-BAC), Rule-based Access Control (Rule-BAC), and Attribute-based Access Control (ABAC). Examples of models that support Mandatory Access Control include Bell-LaPadula and Biba.

mandatory vacation A security policy designed to reduce fraud within an organization. A mandatory vacation policy requires employees to take a vacation outside of the workplace for a specified period, such as at least five consecutive workdays. The goal is to require someone else to perform the employee's job functions, which increases the possibility of exposing any suspicious activities.

mantrap A security control designed to protect against piggybacking. It prevents more than one person from passing through a controlled entry at a time and ensures that each person uses credentials to gain entry.

maximum acceptable outage (MAO) The maximum amount of time a system can be down before critical business functions are affected. MAO is sometimes called *maximum tolerable outage (MTO)* or *maximum tolerable downtime (MTD)*. The MAO helps an organization determine the recovery time objective (RTO).

Media Access Control (MAC) address A 48-bit or 12-digit hexadecimal number (such as 6c-62-6c-ba-73-6c) assigned to a network interface card (NIC). The MAC address uniquely identifies the computer in a network. Also called *physical address* or *hardware address*.

mesh A network topology that provides redundancy with multiple connections. Mesh topologies provide the highest availability of any of the topology methods. See also *bus*, *star*, *token ring*, and *tree*.

Message Digest 5 (MD5) A hashing algorithm used for integrity. MD5 creates a 128-bit hash. The U.S. government considers it cryptographically broken, so it is not recommended for use in many government applications. SHA-3 is a much stronger alternative.

mobile code Code that can be transferred between systems without the user taking explicit action to install it. It can include script code such as JavaScript, compiled code such as Java applets or ActiveX controls, documents with embedded macros, or malware that executes from a USB flash drive when the user inserts the drive into a system.

mobile site An alternative location used in business continuity planning. A mobile site can easily be moved and is useful when an organization doesn't want to designate a specific alternative location. For example, it's possible to set up the inside of a storage container as an alternative location. These are the same storage containers that 18-wheeler tractor-trailers use to haul goods, and they are as easy to move as any other loads hauled over the highways.

multifactor authentication Authentication using more than one factor. There are three factors of authentication: something you know (such as a password), something you have (such as a smart card), and something you are (using biometrics).

near field communication (NFC) Technology in smartphones that allows a user to transfer information to other devices by waving a smartphone close to the other device or tapping the other device. NFC can be used to make purchases with credit cards registered on the phone.

Network Address Translation (NAT) A protocol that translates private IP addresses to public IP addresses and public IP addresses back to private IP addresses. NAT is often installed on proxy servers or routers that are on the edge of the network (between the Internet and the intranet).

Network Discovery Protocol (NDP) Resolves IPv6 addresses to EUI-64 addresses. It is similar to ARP, which resolves IPv4 addresses to MAC addresses. IPv6 uses a 64-bit modified EUI-64 address instead of the MAC address.

Network File System (NFS) Allows computers running different operating systems to access and share files over the network. It allows users running UNIX (and UNIX derivatives) to access files on Microsoft systems, and allows users running Microsoft systems to access files on UNIX-based systems. Sun Microsystems created NFS.

network-based firewall A firewall that provides protection for a network. Traffic to and from the network flows through the network-based firewall. Compare to *host-based firewall*.

network-based IDS (NIDS) An intrusion detection system (IDS) that monitors the network traffic for any type of attack. A NIDS typically has several nodes or agents stationed around the network connected to routers and possibly switches. Each of these nodes monitors the traffic and reports its findings to a NIDS management server.

nonpromiscuous mode A mode used by a sniffer that limits the traffic that the sniffer can capture. The sniffer only captures data sent directly to or from the IP address of the computer running the sniffer. Compare to *promiscuous mode*.

nonrepudiation The practice of ensuring that a party cannot believably deny (or repudiate) taking an action. Nonrepudiation is enforced through audit logging and with digital signatures.

object A resource accessed by a subject. For example, if a user accesses a file, the user is the subject and the file is the object.

one-time password A password used only once. Hardware tokens use synchronous one-time passwords with a hardware device held by the user showing a one-time password as a number or password. The password changes often and is synchronized with a server.

Online Certificate Status Protocol (OCSP) A method of validating certificates with a certificate authority (CA). Clients send the serial number of a certificate to a server known as an OCSP responder. The OCSP responder identifies the health of the certificate with an answer of “good,” “revoked,” or “unknown.”

Open Systems Interconnection (OSI) Model A seven-layer model developed by the International Organization for Standardization (ISO) as a framework for connecting computers and networks together. The seven layers are Physical (layer 1), Data Link (layer 2), Network (layer 3), Transport (layer 4), Session (layer 5), Presentation (layer 6), and Application (layer 7).

packet-filtering firewall A firewall that filters traffic by examining the contents of a packet. A packet-filtering firewall can filter traffic based on IP addresses, subnet addresses, ports, some protocols, or any combination of these.

Password-Based Key Derivation Function 2 (PBKDF2) A key-stretching algorithm that adds a salt of at least 64 bits and then applies a cryptographic function multiple times. Salting helps thwart rainbow table attacks.

penetration test An assessment that starts with a vulnerability assessment. Instead of stopping after discovering vulnerabilities, a penetration test attempts to exploit the discovered vulnerabilities.

personally identifiable information (PII) Information that can be used to personally identify an individual. Many laws mandate the protection of PII.

pharming An attack that redirects users to bogus websites. It manipulates one of the host name resolution methods so that the host name resolves to a different website.

phishing An e-mail-based attack. Attackers send e-mails with the goal of tricking victims into clicking a link or providing sensitive information. The e-mail claims to be from a legitimate company and encourages the user to take an unsafe action.

piggybacking A social engineering tactic. Piggybacking occurs when someone passes through a controlled entry without providing credentials by following closely behind someone who has provided credentials.

plaintext Data that is not in an encrypted format. It is easily readable without any cryptographic or cryptanalysis techniques. Encryption techniques convert plaintext to ciphertext. Also called *cleartext*.

Platform-as-a-Service (PaaS) A type of cloud computing that provides customers with a preconfigured computing platform. Users have access to an operating system and one or more applications hosted on hardware that they can use over the Internet. Users don't have to purchase the hardware and software, and the cloud provider maintains it. Compare to *Infrastructure as a Service (IaaS)* and *Software-as-a-Service (SaaS)*.

Point-to-Point Tunneling Protocol (PPTP) A tunneling protocol used with some virtual private networks (VPNs). RFC 2637 defines it, and Microsoft uses its own version.

policy High-level documents used to provide guidance to members of an organization. A security policy provides direction to employees and is authoritative in nature.

polymorphic virus A virus that has the ability to morph or mutate each time it replicates to another machine or even each time it is run. When the virus mutates, it's more difficult for antivirus (AV) software to detect it unless additional virus signatures are created to recognize the mutated version.

port number Numbers that identify protocols based on assigned well-known ports. For example, SMTP uses TCP port 25. This is not a protocol number, though protocol numbers also identify some protocols. Compare to *protocol number* and *well-known port*.

port scan attack An attempt to detect what ports are open on a system as part of an overall fingerprinting attack. Open ports indicate which services are running on the system

Post Office Protocol version 3 (POP3) A protocol used to receive e-mail. POP3 servers send the e-mail to the client when the client connects. Compare to *Simple Mail Transfer Protocol (SMTP)* and *Internet Message Access Protocol version 4 (IMAP4)*.

primary key A key used in databases that ensures that each row or tuple within a table is unique.

private cloud A cloud-based service available only to users within an organization. As an example, a company can set up cloud-based storage for its employees. Compare to *public cloud*, *community cloud*, and *hybrid cloud*.

private IP address An IP address used within an organization's private network. RFC 1918 formally defines private IP addresses. Compare to *public IP address*.

private key A key used with public key cryptography for asymmetric encryption. A private key is part of a matched pair (matched to a public key). A private key is kept private and not shared with other entities. Compare to *public key* and *session key*.

procedure Specific action steps to accomplish tasks. Procedures (and guidelines) are derived from policies.

promiscuous mode A mode used by a sniffer that increases the amount of traffic that the sniffer can capture. The sniffer can capture all data that reaches the sniffer regardless of the destination IP address. Compare to *nonpromiscuous mode*.

protected health information (PHI) Personally identifiable information that includes an individual's medical and health history.

protocol number A number embedded in the first octet of TCP and UDP headers that identifies some protocols. For example, ICMP uses protocol number 1 and IGMP uses protocol number 2. This is not a port number, though port numbers also identify some protocols. Compare to *port number*.

proxy server A server used as an intermediary, or proxy, for internal clients accessing Internet resources. The proxy server retrieves the web page on the client's behalf and returns it to the client.

public cloud A cloud-based service provided by a third-party vendor and available to anyone. As an example, Apple provides cloud-based storage via its iCloud service. Compare to *private cloud*, *community cloud*, and *hybrid cloud*.

public IP address An IP address that is used on the public Internet. Compare to *private IP address*.

public key A key used with public key cryptography for asymmetric encryption. A public key is part of a matched pair (matched to a private key). A public key is freely shared and distributed within a certificate. Compare to *private key* and *session key*.

public key infrastructure (PKI) All the components necessary to create, manage, distribute, validate, and revoke certificates. A PKI is based on the X.509 protocol. The X.509 standard identifies many of the components and formats used by the PKI, certificate authorities (CAs), and certificates.

qualitative analysis A risk assessment method that uses subjective opinions from experts to identify the impact of a risk. It often categorizes a risk using words such as "low," "medium," and "high." Compare with *quantitative analysis*.

quantitative analysis A risk assessment method that uses numerical-based data such as monetary figures to identify the actual cost associated with a risk. Compare with *qualitative analysis*.

radio-frequency identification (RFID) A system used for identification, tracking, asset management, and inventory control. Products are "tagged" with an RFID tag, which is a small electronic device that marks the product. RFID readers can then read the tag to get information about the product.

rainbow table A table of passwords using a predefined character set and the hash for each of the passwords. Attackers use rainbow tables to perform rainbow table lookups and identify passwords from the hash.

ransomware A type of Trojan malware that takes control of a user's computer and then demands a ransom from the user to get control back.

RC4 A symmetric encryption standard. This standard is sometimes called Ron's Code or Rivest's Cipher after its inventor, Ron Rivest. Many experts have speculated that RC4 has been cracked, and it is no longer recommended for use.

recovery point objective (RPO) A recovery term associated with backups and databases that identifies the amount of data (in terms of time such as hours or days) that is acceptable to lose if a failure occurs. The RPO dictates the amount of resources needed to protect data to prevent loss in the case of a failure.

recovery time objective (RTO) The maximum amount of time that can be taken to restore a system or process to operation. If a failure takes longer than the RTO to restore, then it impacts the mission. The maximum acceptable outage (MAO) helps an organization determine the RTO.

Remote Authentication Dial-In User Service (RADIUS) A service that provides centralized authentication, authorization, and accounting (AAA) for remote clients. Remote users are authenticated against a credential's database, their account is checked to verify that they are authorized to use remote access, and their remote access activity is tracked in an accounting log.

residual risk The amount of risk that remains after steps have been taken to mitigate risk. In other words, residual risk = total risk – mitigated risk. Management is responsible for any losses that occur as a result of residual risk.

Reverse Address Resolution Protocol (RARP) You can think of this as a simplified version of DHCP. A client has a MAC address and uses RARP to get an IP address. It provides the clients some TCP/IP configuration information, but doesn't have the full capabilities of DHCP.

risk The probability or likelihood of a threat exploiting a vulnerability, resulting in a loss. A threat is any activity that can be a possible danger. A vulnerability is a weakness, and a loss represents a negative event for an organization. The level of loss represents the impact.

risk assessment A point-in-time evaluation of potential risks. It looks at the current situation and then attempts to determine what risks exist and how to address them.

Role-based Access Control (Role-BAC) model An access control model that uses roles to determine access. Subjects (such as users) are placed into roles, and access to objects (such as files) is granted to the roles. Other access control models are Mandatory Access Control (MAC), Discretionary Access Control (DAC), Rule-based Access Control (Rule-BAC), and Attribute-based Access Control (ABAC).

rootkit A set of programs that can run on a system largely undetected. It has root-level access to the system, similar to how a root-level administrator has full and complete control over a system.

RSA An algorithm used for encryption and decryption in public key cryptography. RSA uses large prime numbers to create secure, matching public and private keys. Its strength lies in the fact that it's computationally infeasible to factor the composite number created from these prime numbers as long as the prime numbers are sufficiently large.

Rule-based Access Control (RBAC) model An access control model that uses rules to determine access. As an example, routers use rules within an access control list (ACL) to define traffic allowed in or out of a network. Other access control models are Mandatory Access Control (MAC), Discretionary Access Control (DAC), Role-based Access Control (Role-BAC), and Attribute-based Access Control (ABAC).

Safe Harbor A program that helps U.S. organizations comply with the requirements of the Data Protection Directive. It includes seven principles that organizations agree to follow.

safeguard A means, method, action, technique, process, procedure, or device that reduces the vulnerability of a system or the possibility of a threat exploiting a vulnerability in a system. Controls are risk management tools. The terms *control*, *countermeasure*, and *safeguard* are often used interchangeably.

salting A process that adds random bits to a password. Salting helps thwart rainbow table attacks and is used by bcrypt and Password-Based Key Derivation Function 2 (PBKDF2).

scareware Malware that describes itself as free antivirus software to trick users into downloading it. It typically attempts to scare the user with a pop-up or other message that indicates the user's system is infected with a virus. Sometimes called *rogueware*.

Secure File Transfer Protocol (SFTP) A protocol that uses Secure Shell (SSH) to encrypt FTP traffic. SFTP functions similarly to FTP, but the encryption prevents sniffing attacks from capturing data in cleartext. Compare to *File Transfer Protocol (FTP)* and *Trivial FTP (TFTP)*.

Secure Hashing Algorithm 1 (SHA-1) A hashing algorithm used for integrity. SHA-1 creates a 160-bit hash. Due to vulnerabilities, it is no longer recommended for use.

Secure Hashing Algorithm 2 (SHA-2) A hashing algorithm used for integrity. SHA-2 is a hash of 224, 256, 384, or 512 bits. It is an improvement to SHA-1, though many experts believe it will be broken because it functions similarly to SHA-1.

Secure Hashing Algorithm 3 (SHA-3) The winner of the NIST hash function competition from 2007 to 2012. It is not based on SHA-1 or SHA-2, but instead is a subset of the Keccak algorithm. It creates hashes of 224, 256, 384, or 512 bits.

Secure/Multipurpose Internet Mail Extensions (S/MIME) The standard used to encrypt and digitally sign e-mail. Almost all e-mail applications that support encryption and digital signatures use S/MIME.

Secure Real-time Transport Protocol (SRTP) A security protocol that provides confidentiality, message authentication, and replay protection for audio and video traffic, including Voice over Internet Protocol (VoIP).

Secure Shell (SSH) A basic encryption protocol used to create a secure session between two computers. SSH is used to encrypt other protocol traffic such as File Transfer Protocol (FTP) (called Secure FTP [SFTP]) and Secure Copy (SCP) traffic. It is a secure alternative to tools such as Telnet, rlogin, rsh, and rexec. SSH uses port 22.

Secure Sockets Layer (SSL) An older encryption protocol used to encrypt different types of traffic. SSL is susceptible to a POODLE (Padding Oracle On Downgraded Legacy Encryption) attack. Most organizations have stopped using SSL in favor of the designated replacement, Transport Layer Security (TLS). SSL commonly uses TCP port 443.

security policy A written document that provides the organization with a high-level view of its security goals. A security policy is authoritative in nature and provides direction for the creation of guidelines and procedures.

separation of duties A security principle that ensures that no single person has complete control over a process. When properly implemented, it can significantly reduce the risk of fraud within an organization.

session key A key used with symmetric encryption. It is sometimes called a symmetric key. Compare to *private key*, *public key*, and *asymmetric key*.

shoulder surfing The practice of attempting to gain information by looking over a user's shoulder as he or she is entering data.

signature-based detection A method of virus detection used to detect known viruses. Viruses have specific characteristics that can be used to identify them uniquely. The signature can be a unique characteristic such as a specific byte pattern within the virus. Intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) can also use signatures to detect known attack methods.

Simple Mail Transfer Protocol (SMTP) A protocol used to send e-mail. SMTP sends e-mail from e-mail clients to an e-mail server. E-mail servers use SMTP to send and receive e-mail between other SMTP servers. SMTP uses TCP port 25. Compare to *Post Office Protocol version 3 (POP3)* and *Internet Message Access Protocol version 4 (IMAP4)*.

Simple Network Management Protocol (SNMP) A protocol used to manage network devices (such as routers and layer 3 switches) on a network. SNMP agents on the network devices send traps (error or event notifications) and other notification data back to a central server running SNMP management software. SNMP agents receive data on UDP 161 and send traps and notifications on UDP port 162.

single sign-on (SSO) The practice of allowing a subject (such as a user) to authenticate once and use the same credentials to access additional resources without authenticating again. This increases security because users have to remember only a single set of credentials and are less likely to write their credentials down. Federated access SSO systems allow users to access systems owned and managed by different organizations by logging on once using credentials recognized by the federated access system.

sniffer A protocol analyzer or packet sniffer. It is capable of capturing and analyzing packets transmitted over a network. A popular sniffer is Wireshark.

social engineering The practice of using primarily nontechnical means to get people to give up sensitive data or to perform actions they wouldn't normally perform. A social engineer uses deception and fraud to trick or manipulate unsuspecting users.

Software-as-a-Service (SaaS) A type of cloud computing that provides one or more applications to users. Users access the applications with a web browser over the Internet. Web-based e-mail is an example of SaaS used by many people. Compare to *Infrastructure-as-a-Service (IaaS)* and *Platform-as-a-Service (PaaS)*.

something you are One of the three factors of authentication. This factor includes the use of biometrics to authenticate an individual based on fingerprints, retina, or other facial characteristics, keystroke dynamics, and handwriting.

something you have One of the three factors of authentication. This factor includes items such as proximity cards, smart cards, hardware tokens, and identification badges.

something you know One of the three factors of authentication. This factor includes knowledge such as passwords, personal identification numbers, mother's maiden name, or even personal information such as the name of your first pet.

spam Unsolicited commercial e-mail (UCE) and other unwanted electronic messages sent in bulk. Spam is most often associated with e-mail, but can also be delivered via other means, such as instant messaging. It usually includes advertisements, but can also be used in different types of attacks launched from botnets and phishing.

spam filter An e-mail scanner that attempts to detect and block unsolicited e-mail (spam). Spam filters can be installed on end-user systems and e-mail servers and at the boundary of the network to scan all traffic from the Internet.

spyware Software that is installed on a user's system without the user's knowledge or consent. Spyware is considered a form of malware because it can cause damage to an individual.

SQL injection An attack that injects SQL code to read and manipulate databases. SQL injection attacks are mitigated with input validation techniques and stored procedures.

standard A proven norm or method. Standards are typically external to an organization but can influence the organization's policies, guidelines, and procedures.

star Network configuration where all computing devices are connected to a central device, such as a hub or a switch. Most networks use a star configuration, with all of the devices connected to the central device with twisted pair cable. See also *bus*, *mesh*, *token ring*, and *tree*.

stateful inspection firewall A firewall that filters traffic based on the state of existing connections. It identifies active connections as they are created and monitors the status of these connections in a state table within the firewall.

steganography The practice of hiding data within data. When used to hide data within files, it modifies the least significant bits of bytes within the files.

subject An entity that can access a resource referred to as an object. For example, if a user accesses a file, the user is the subject and the file is the object.

symmetric encryption A process of encrypting and decrypting data using a single key. Symmetric encryption algorithms are constant, while the keys used to encrypt different sets of data are changed to prevent compromise. A popular symmetric encryption is Advanced Encryption Standard (AES).

TCP/IP Model A model developed by the Defense Advanced Research Projects Agency (DARPA). Some references list it with four layers (Application, Transport, Internet, and Link). Other references list it with five layers (Application, Transport, Network, Data Link, and Physical).

Telnet A basic protocol used to interact with a remote system using text-based commands. Telnet sends data (including login credentials) across the network in cleartext, which attackers can capture and read using a sniffer. Secure Shell (SSH) is a secure alternative.

Terminal Access Controller Access Control System + (TACACS+) An alternative to RADIUS to provide centralized authentication, authorization, and accounting (AAA) services for remote clients. TACACS+ provides improvements over TACACS, and both TACACS and TACACS+ use port 49. TACACS+ uses Transmission Control Protocol (TCP). TACACS uses User Datagram Protocol (UDP).

threat Any activity that can be a possible danger. When a threat exploits a vulnerability, an organization can suffer losses.

token ring A network configuration where all computing devices are connected in a logical circle. Devices can transmit data onto the network only when they have a logical token that is passed from device to device. See also *bus*, *mesh*, *star*, and *tree*.

Transmission Control Protocol (TCP) A connection-oriented protocol that provides guaranteed reliable communication for devices on a network. A three-way handshake establishes a TCP connection. TCP uses packet sequencing, and the destination acknowledges every packet that it receives.

Transport Layer Security (TLS) The designated replacement for Secure Sockets Layer (SSL). It is based on SSL and is formally defined in RFC 5246. Almost all HTTPS sessions today are encrypted with TLS due to known vulnerabilities with SSL. TLS commonly uses TCP port 443.

trapdoor Code embedded in an application used to provide access to the application, the application's code, or its data via a covert method. It is also called a *backdoor*.

tree A combination of a bus topology and a star topology. Instead of connecting multiple computer nodes in the bus configuration, it connects multiple star networks along a type of a bus network. See also *bus*, *mesh*, *star*, and *token ring*.

Trivial FTP (TFTP) A stripped-down version of FTP that provides very basic capabilities. It doesn't support authentication and doesn't use TCP to establish a session. However, it is very useful when transferring configuration files to and from network devices. TFTP uses UDP port 69.

Trojan horse Malware that looks like one thing but is actually something different. Users are often tricked into installing the malware, thinking that it will provide some benefit. When the application is installed, the Trojan horse also installs the malware.

tuple A row storing data within a database table. Some database vendors call tuples *rows*, while other vendors call them tuples. Similarly, some vendors refer to database table columns as *attributes*.

User Datagram Protocol (UDP) A connectionless protocol that uses a best effort to send data without verification. Instead of checking to see whether a connection exists with another system before sending data, it simply sends it. In comparison, Transport Control Protocol (TCP) is a connection-oriented protocol that ensures a connection exists before sending data. As an example, Trivial File Transport Protocol (TFTP) uses UDP, but File Transport Protocol (FTP) uses TCP.

virtual local area network (VLAN) A segmented LAN created on a switch. A VLAN allows an organization to segment traffic with a lot more flexibility when compared to segmenting traffic with a router.

virtual private network (VPN) A connection that provides access to a private network over a public network such as the Internet. VPNs use tunneling protocols (such as Point-to-Point Tunneling Protocol [PPTP] or Layer 2 Tunneling Protocol [L2TP]) to secure the traffic.

virus A form of malicious software (malware). It is an application or a piece of code that causes unexpected and usually negative events on computers. One of the key characteristics of a virus is that the infected file must be executed for the virus to run.

Voice over Internet Protocol (VoIP) The different technologies used to transmit multimedia and voice communications over IP networks, including the Internet.

volatile RAM Any memory that requires power to hold the data. When a system is powered down, any data (including potential forensic evidence) within volatile RAM is lost.

vulnerability A weakness. It can be a weakness in a system, a configuration, a process, hardware, software, or any other aspect of a system. If a threat can exploit a vulnerability, an organization can suffer losses.

vulnerability assessment A process used to discover vulnerabilities. Vulnerability assessments can use technical tools to scan networks and systems for vulnerabilities. They can also include nontechnical means such as social engineering tactics to determine whether employees are susceptible to social engineering attacks.

warm site An alternative location used in business continuity planning. A warm site is a compromise between a cold site and a hot site. The organization makes compromises with costs and time. Compare to *hot site*, *cold site*, and *mobile site*.

well-known port Logical port numbers from 0 to 1,023 mapped to specific protocols by the Internet Assigned Numbers Authority (IANA). For example, IANA assigned port 80 to the HTTP protocol. Port numbers are assigned as Transmission Control Protocol (TCP) ports and User Datagram Protocol (UDP) ports.

whitelist A list of authorized applications. Application whitelisting identifies specific applications that can run on a system and all other applications are blocked. Compare to *blacklist*.

Wi-Fi Protected Access (WPA/WPA2) Wi-Fi–related security protocols. WPA was introduced as a replacement for Wired Equivalent Privacy (WEP) to secure wireless networks. WPA2 was created as a permanent replacement and is formalized as IEEE 802.11i. WPA and WPA2 can operate in either Personal mode or Enterprise mode. Personal mode uses a preshared key (PSK), and Enterprise mode uses an authentication server such as a Remote Dial-in User Authentication Service (RADIUS) or 802.1x server.

WiMAX Worldwide Interoperability for Microwave Access, the commercial name for 802.16-related products. WiMAX provides broadband wireless access over a large metropolitan area, such as a city, with speeds comparable to wired broadband access.

Wired Equivalent Privacy (WEP) The first security algorithm used for 802.11 wireless networks. WEP has several security issues and has been deprecated in favor of WPA and WPA2. It should not be used.

worm Malware that travels over the network looking for systems to infect. Worms do not require any type of interaction to execute.

zero day exploit An attack that takes advantage of unpublished vulnerabilities. In some cases, the vendor knows about the exploit but has not released a patch yet.

This page intentionally left blank

INDEX

* Integrity Axiom, 54

*nix logs, 350–351

3DES, 478

3G, 104–105

4G, 105

A

AAAs of security, 17–18

acceptable use policy, 197, 316, 346, 408, 411

access control entries (ACEs), 48

access control lists (ACLs), 70, 125

access control matrix, vs. capability tables, 55

access control systems

Attribute-based Access Control (ABAC), 51–52

comparing access control models, 47–55

comparing identification, authentication, and authorization, 27–28

comparing subjects and objects, 44–45

Discretionary Access Control (DAC) model, 47–49

implementing, 44–47

logical access controls, 47

Mandatory Access Control (MAC) model, 52–55

non-Discretionary Access Control (non-DAC) models, 49

Role-based Access Control (Role-BAC) model, 49–50

Rule-based Access Control (Rule-BAC) model, 51

accidental threat sources, 246

accidental threats, 162

account de-provisioning, 58

account lockout policies, 57, 347, 411

accountability, 18

accounting, 18

accreditation. *See* certification and accreditation

ACPO Good Practice Guide for Digital Evidence, 447

Address Resolution Protocol. *See* ARP

administrators, 375

Advanced Encryption Standard. *See* AES

advanced persistent threats (APTs), 159–160

vs. insider attacks, 161

adversarial threat sources, 245–246

adverse impacts, 245

AES, 100, 122, 474, 477

AH, 88, 89, 492

AIC security triad, 9

AirMagnet WiFi Analyzer Pro, 191

Amazon Elastic Compute Cloud (Amazon EC2), 142

Android, 321, 323–324

annual loss expectancy (ALE), 255–256

annual rate of occurrence (ARO), 255–256

anti-analysis malware, 443

antivirus software, 224–225, 304, 322

See also viruses

Apple

App Store, 323

backdoors, 324

appliances, 141, 143

content-filtering, 229–230

application firewalls, 127–128

application review, as a countermeasure, 175

application whitelisting and blacklisting, 234, 323

APT1, 288

APT1: Exposing One of China's Cyber Espionage Units, 288, 453

APTs. *See* advanced persistent threats (APTs)

architecture, internetwork trust architectures, 91–97

archiving data, 377

armored viruses, 209

See also viruses

ARP, 81–82

ASCII, 72

asset management, 390

data, 391

hardware, 390–391

software, 391

assets, identifying, 252

Asymmetric DSL (ADSL), 121

asymmetric encryption, 479–480

Diffie-Hellman, 483, 484

elliptic curve cryptography (ECC), 484

Elliptic Curve Digital Signature Algorithm (ECDSA), 484

protecting e-mail with S/MIME, 484–490

RSA, 481

Secure Shell (SSH), 484

Secure Sockets Layer (SSL), 483

Transport Layer Security (TLS), 481–483

asynchronous dynamic passwords, 34

attackers

accidental threats, 162

advanced persistent threats (APTs), 159–160

hackers and crackers, 158

insider attacks, 160–161

overview, 157

phreaks, 162

script kiddies, 161–162

white hats, black hats, and grey hats, 158–159

attacks and countermeasures

application review, 175

botnets and zombies, 165–167

brute-force attacks, 182

buffer overflow attacks, 177

- attacks and countermeasures (*Cont.*)
 - chosen-plaintext attacks, 502
 - ciphertext-only attacks, 502
 - code signing, 175–176
 - covert channel, 190
 - cross-site request forgery, 180–182
 - cross-site scripting, 180
 - cryptanalysis attacks, 501–502
 - data inference attacks, 386
 - denial of service (DoS) attacks, 164–165, 263
 - dictionary attacks, 182–183
 - distributed denial of service (DDoS) attack, 165, 263
 - evil twins, 191–192
 - fraggle attacks, 173–174
 - hardening systems, 163
 - increasing user awareness, 163
 - injection attacks, 178–180
 - input validation, 174–175
 - known-plaintext attacks, 502
 - man-in-the-middle attacks, 171–172
 - password attacks, 182–185
 - patching systems, 163
 - phishing, 185–187
 - ping sweeps, 170
 - port scans, 170
 - replay attacks, 173
 - rogue access points, 191
 - salami attacks, 171
 - sandboxing, 176–177
 - session hijacking, 172–173
 - smishing, 188
 - smurf attacks, 173–174
 - sniffing attacks, 167–170
 - social engineering, 183
 - software security, 174
 - spam, 185
 - spear phishing and whaling, 187–188
 - spoofing, 163
 - vishing, 188
 - wardriving, 192
 - WIDS and WIPS as countermeasures, 191
 - wireless attacks and countermeasures, 190–191
 - WPA cracking attacks, 192
 - zero day exploits, 188–190
 - Attribute-based Access Control (ABAC), 51–52
 - audit logging, 12, 16, 344–346, 348
 - *nix logs, 350–351
 - auth log, 351
 - firewall logs, 352
 - maillog, 351
 - managing audit logs, 353–354
 - operating system logs, 348–349
 - proxy server logs, 351–352
 - reviewing logs, 352–353
 - storing logs on remote systems, 349–350
 - sulog, 350
 - syslog, 350
 - audit passwords, 31
 - auditing, 148
 - and account lockout, 347
 - audit trails, 348
 - clipping levels, 346–347
 - ISACA, 356
 - minimum requirements for federal agencies, 345
 - passwords, 355
 - PCI DSS compliance audits, 357–358
 - physical access controls, 358
 - purpose of, 346
 - reviewing logs, 352–353
 - security policies, 355
 - in security policies, 410
 - through an inspection process, 343
 - through logs, 343, 344–346
 - See also* audit logging; security audits
 - auth log, 351
 - authentication, 17, 27–28, 485
 - centralized vs. decentralized, 42–43
 - certificates, 493
 - device, 43–44
 - multifactor, 37–38
 - offline, 43
 - overview, 28
 - remote access, 134
 - reviewing identification, 38
 - and security policies, 411
 - single sign-on, 38
 - something you are, 35–37
 - something you have, 33–35
 - something you know, 29–33
 - three factors of, 29
 - authentication header. *See* AH
 - authenticity, 466
 - authorization, 18, 27–28
 - provisioning and, 56–57
 - autonomy-based detection, 282
 - See also* intrusion prevention systems (IPSs)
 - AV software. *See* antivirus software
 - availability, 12–13
- ## B
- backdoors, 217–218, 324
 - backups, 12
 - backup plans in security policies, 410
 - backup policy, 316
 - full, 332
 - full/differential backup strategy, 333–334
 - full/incremental backup strategy, 332–333
 - overview, 331–332
 - Bcrypt, 478
 - behavioral biometrics, 36
 - Bell, David Elliott, 53–54
 - Bell-LaPadula model, 53–54
 - Berkeley Internet Name Domain (BIND), 83
 - Biba model, 54
 - Big Data, securing, 387
 - biometrics, 35–37, 38
 - errors, 37
 - BitLocker, 324
 - bits, 67
 - black box testing, 289
 - black hats, 158–159

- blacklists, 234
 - block ciphers, 476
 - Blowfish, 478
 - bluebugging, 104
 - bluejacking, 104
 - bluesnarfing, 104
 - Bluetooth, 104
 - Bogachev, Evgeniy Mikhailovich, 214
 - boot sectors, 209
 - See also* viruses
 - botnets, 165–167
 - Brewer, David F.C., 55
 - Brewer-Nash model, 55
 - brute-force attacks, 182
 - buffer overflow attacks, 177
 - bus topology, 75
 - business continuity plans (BCPs), 317, 417–419
 - alternate locations, 425–427
 - business impact analysis (BIA), 419–422
 - vs. disaster recovery plans, 423–424
 - restoration planning, 424
 - and security policies, 410
 - testing and drills, 424–425
 - business impact analysis (BIA), 419–420
 - maximum acceptable outage, 420
 - output, 421–422
 - recovery point objective, 421
 - recovery time objective (RTO), 420
 - BYOD, 320–321
- ## C
- cable modems, 121
 - CA/Browser Forum, 176
 - Candidate Information Bulletin (CIB), 5
 - canvas fingerprinting, 218
 - capability tables, vs. access control matrix, 55
 - CBA. *See* cost-benefit analyses (CBA)
 - CBK. *See* Common Body of Knowledge (CBK)
 - CCMP, 99, 100
 - cellular connections, 121
 - centralized authentication, 42–43
 - CERT Division, 430
 - certificate authorities (CAs), 135, 495–496
 - Trusted Root Certification Authorities, 497–498
 - certificates, 492–495
 - alternative certificate trusts, 500–501
 - revoking, 498
 - trust chain, 497, 500–501
 - validating, 498–499
 - certification and accreditation, 392
 - Common Criteria (CC), 393
 - Evaluation Assurance Levels (EALs), 393
 - and security assessments, 392–393
 - TCSEC, 394
 - using a risk management framework, 394–395
 - Certified Information Systems Auditor (CISA)
 - certification, 356
 - Certified Information Systems Manager (CISM)
 - certification, 356
 - Certified Information Systems Security Professional
 - certification. *See* CISSP certification
 - chain-of-custody forms, 446
 - Challenge Handshake Authentication Protocol. *See* CHAP
 - change authorization board (CAB), 362
 - change control, 317–318, 361–363
 - change management, 361–363
 - change review board (CRB), 362
 - CHAP, 135, 173
 - Children's Online Privacy Protection Act (COPPA), 456
 - Chinese Wall model, 55
 - chosen-plaintext attacks, 502
 - CIA security triad, 9
 - CIFS, 141
 - ciphertext, 466
 - ciphertext-only attacks, 502
 - CISA certification, 356
 - CISM certification, 356
 - CISSP certification, 5
 - Drag & Drop questions, 5, 6
 - Hotspot questions, 5, 7
 - Clark, David, 54–55
 - Clark-Wilson model, 54–55
 - classification of resources, 410
 - clipping levels, 346–347
 - cloud computing
 - cloud operation models, 145–146
 - community cloud, 146
 - compliance, 148
 - data control and third-party outsourcing, 147–148
 - data management policies, 379
 - hybrid cloud, 146
 - overview, 144–145
 - privacy, 147
 - private cloud, 146
 - public cloud, 145
 - storage, 146
 - Cloud Computing Security Requirements Guide*, 144, 379
 - coaxial cable, 68
 - COBIT framework, 356
 - Code of Ethics, 2–3, 414–415
 - See also* security policies
 - code signing
 - certificates, 176, 493
 - as a countermeasure, 175–176
 - cognitive passwords, 29–30, 56
 - cold sites, 13, 427
 - command injection, 179–180
 - Common Access Cards (CACs), 38
 - Common Body of Knowledge (CBK), 3, 5
 - Common Criteria (CC), 393
 - Common Internet File System. *See* CIFS
 - Common Vulnerabilities and Exposures (CVE), 235, 254
 - compensating controls, 309–310
 - compliance, PCI DSS compliance audits and reports, 357–358
 - computer abuse, 448–449
 - computer crime, 448–449
 - fraud and embezzlement crime, 450–452
 - Computer Emergency Response Team (CERT), 210

- computer forensics, 439
 - acquiring evidence, 444–446
 - analyzing evidence, 447
 - authenticating evidence, 446–447
 - chain-of-custody forms, 446
 - evidence guidelines and principles, 447–448
 - incident handling, 439–442
 - preserving the scene, 442–443
 - three phases of investigations, 443–447
 - toolkits, 445
 - computer incident response team (CIRT), 264–265
 - Computer Security Incident Handling Guide*, 262
 - confidentiality, 9–10, 465
 - See also* cryptography
 - confidentiality, integrity, and availability (CIA), 245
 - See also* availability; confidentiality; integrity
 - ConfigMgr, 319, 322
 - configuration management, 317–318
 - Federal Desktop Core Configuration (FDCC), 360
 - overview, 358–359
 - system baselines, 359
 - United States Government Configuration Baseline (USGCB), 359–360
 - using Group Policy for, 360–361
 - using imaging for, 359–360
 - connectionless, 71
 - connection-oriented, 71
 - content-filtering appliances, 229–230
 - continuing professional education credits. *See* CPE credits
 - control plane, 142
 - converged communications, 123
 - cookies, 173
 - corrective controls, 309
 - cost of control, 255–256
 - cost-benefit analyses (CBA), 304–305
 - countermeasures. *See* attacks and countermeasures; security controls
 - covert channel, 190
 - CPE credits, 7–8
 - crackers, 158
 - credential management systems, 31
 - Credential Manager tool, 31
 - credit card verification value (CVV), 56
 - criminal history and background, 3
 - Crossover Error Rate (CER), 37
 - cross-site request forgery, 180–182
 - cross-site scripting, 180
 - cryptanalysis, 467
 - attacks, 501–502
 - cryptographic systems, 467
 - cryptography
 - asymmetric encryption, 467
 - authenticity, 466
 - basic concepts, 465–466
 - ciphertext, 466
 - confidentiality, 465
 - data sensitivity, 467–468
 - defined, 467
 - encryption, 466
 - encryption algorithm, 467
 - hashing, 467
 - integrity, 465
 - managing cryptographic keys, 501
 - plaintext, 466
 - regulatory requirements, 468
 - symmetric encryption, 467
 - CryptoLocker, 161, 213–214, 218, 222, 227
 - installing, 214
 - CryptoWall, 222
 - CSMA/CA, 75
 - CSMA/CD, 74
 - CSRF. *See* cross-site request forgery
 - Current Activity Updates, 227
 - custodians, 375
 - Cyber Awareness Alerts, 227
 - Cyber Awareness Bulletins, 227
 - Cyber Awareness Tips, 227
 - cyberbullying, 449
 - cyberstalking, 449
 - cyberwarfare, 449
- ## D
- data
 - asset management, 391
 - classifying, 371–374
 - marking and labeling, 374
 - persistent data, 445
 - protecting from cradle to grave, 375
 - roles and responsibilities for, 374–375
 - sensitivity, 467–468
 - volatile data, 445
 - data at rest, 375–376
 - Data Breach Investigations Report*, 160
 - data breaches
 - defining, 453
 - Epsilon data breach, 389
 - data classification, 371–374, 409
 - Data Definition Language (DDL), 385
 - data destruction, 147–148
 - data diddling, 386
 - data in motion, 375–376, 468
 - data in use, 376
 - data inference, 386, 455
 - data leakage, 143, 322
 - data loss prevention (DLP), 380
 - data management policies
 - archiving and retention requirements, 377
 - cloud computing, 379
 - deduplication, 379–380
 - information rights management, 381–382
 - regulatory requirements, 387–390
 - removing data remnants, 377–379
 - securing Big Data, 387
 - social network usage, 381
 - storage media, 376–377
 - transmission, 377
 - See also* databases
 - Data Manipulation Language (DML), 385
 - data normalization, 384
 - data owners, 374
 - data plane, 142
 - data portability, 147
 - data remnants, 377–379

- data resilience, 148
 - data transmissions
 - data management policies, 377
 - securing, 99
 - databases, 382–383
 - communicating with, 385
 - foreign keys, 384
 - OLTP vs. OLAP, 385–386
 - primary keys, 384
 - relationships, 383–384
 - tuples, 384
 - views, 384–385
 - datagrams, 67
 - DDoS. *See* distributed denial of service (DDoS) attack
 - decentralized authentication, 42–43
 - deduplication, 379–380
 - defense diversity, 128–129
 - defense in depth, 15–16, 279
 - degaussing, 378
 - denial of service (DoS) attacks, 164–165, 263
 - de-provisioning, 58
 - DES, 478
 - detective controls, 308
 - deterrent controls, 310
 - device authentication, 43–44
 - device fingerprinting, 43–44
 - device security, 320, 325
 - application whitelisting and blacklisting, 323
 - bringing your own device (BYOD), 320–321
 - endpoint encryption, 323–324
 - mobile device management, 322
 - telecommuting, 321
 - thin clients, 322–323
 - Trusted Platform Module (TPM), 324
 - USB devices, 322
 - using corporately owned devices, 321
 - virtualization, 323
 - DHCP, 80–81, 97
 - Diameter, 137
 - dictionary attacks, 182–183
 - Diffie-Hellman, 483, 484
 - digital rights management, 381
 - digital signatures, 17, 485–487
 - Digital Subscriber Line. *See* DSL
 - direct losses, vs. indirect losses, 413
 - directive controls, 310
 - disaster recovery plans (DRPs), 317, 417–419, 422–423
 - vs. business continuity plans, 423–424
 - testing and drills, 424–425
 - Discretionary Access Control (DAC) model, 47–49
 - discretionary access control lists (DACLS), 48
 - distributed denial of service (DDoS) attack, 165, 263
 - DLP. *See* data loss prevention (DLP)
 - DMZ, 92, 93, 128–129
 - and firewall logs, 348
 - DNS, 82–83, 145
 - Domain Name System. *See* DNS
 - DoS. *See* denial of service (DoS) attacks
 - Drag & Drop questions, 5, 6
 - drive-by downloads, 187, 221–222
 - Dropbox, 147
 - DSL, 121
 - Duchak, Douglas, 161
 - due care, 19–20
 - due diligence, 19
 - dumpster diving, 195
 - Dynamic Host Configuration Protocol. *See* DHCP
 - dynamic passwords, 29, 34
 - dynamic SQL attacks, 178
- ## E
- EAP, 135–136
 - EAP-TLS, 136
 - EAP-Tunneled TLS (EAP-TTLS), 135–136
 - Ebbers, Bernard, 450
 - ECDH. *See* Diffie-Hellman
 - e-discovery, 146
 - electromagnetic interference (EMI), 69
 - electronic discovery. *See* e-discovery
 - elliptic curve cryptography (ECC), 484
 - Elliptic Curve Digital Signature Algorithm (ECDSA), 484
 - e-mail, 197
 - delivering malware, 223
 - encrypting, 488–490
 - protecting with S/MIME, 484–490
 - e-mail protocols, 87–88
 - e-mails, certificates, 493
 - embezzlement, 450–452
 - emergency response plans, 423
 - See also* business continuity plans (BCPs); disaster recovery plans (DRPs)
 - Encapsulated RSPAN (ERSPAN), 120
 - Encapsulating Security Protocol. *See* ESP
 - encryption, 9–10, 106, 147, 375–376, 466
 - asymmetric encryption, 467, 479–490
 - certificates, 493
 - of e-mail, 488–490
 - encryption algorithm, 467
 - endpoint encryption, 323–324
 - IPsec, 492
 - one-way encryption with hashing algorithms, 469–472
 - Pretty Good Privacy (PGP), 490
 - steganography, 490–491
 - symmetric encryption, 467, 474–479
 - endorsement keys, 324
 - enforcement, of security policies, 411, 412
 - entitlement, 57–58
 - environmental threat sources, 246
 - E-Privacy Directive, 454
 - Epsilon data breach, 389
 - ESP, 88–89, 492
 - Ethernet, 73–74
 - ethics statements, 411
 - European Union, privacy directives, 454–455
 - evil twins, 191–192
 - exam. *See* SSCP certification exam
 - Exposing One of China's Cyber Espionage Units*, 159
 - exposure factor, 255
 - Extended Binary Coded Decimal Interchange Code (EBCDIC), 72
 - Extended Validation (EV), 176

eXtensible Access Control Markup Language. *See* XACML
 Extensible Authentication Protocol. *See* EAP
 external testing, 289
 extranets, 92

F

Faceniff, 173
 failover clusters, 329–330
 False Acceptance Rate (FAR), 37
 False Rejection Rate (FRR), 36
 fault tolerance
 for disks, 325–326
 failover clusters, 329–330
 overview, 325
 redundant connections, 330–331
 See also RAID
 FC, 141
 Federal Desktop Core Configuration (FDCC), 360
 federated access, 40–41
 See also single sign-on authentication
 feedback loops, 441–442
 fiber-optic cabling, 69
 Fibre Channel. *See* FC
 Fibre Channel over Ethernet (FCoE), 123
 file integrity checkers, 284–285
 File Transfer Protocol. *See* FTP
 fingerprinting the system, 290, 295
 fingerprints, 35
 FireEye, 213, 214
 Firesheep, 173
 firewalls
 application, 127–128
 basic, 98
 comparing network-based and host-based, 129–130
 defense diversity, 128–129
 logs, 352
 next-generation, 128
 overview, 125
 packet-filtering, 125–127
 stateful inspection, 127
 first responders, 442–443
 forensics. *See* computer forensics
 Fox-IT, 214
 fraggle attacks, 173–174
 frames, 67, 70
 fraud, 450–452
 FTP, 71, 84–85
 full knowledge testing. *See* white box testing
 functional exercises, 425

G

GameOver Zeus, 161, 166, 217–218, 227
 See also malware
 Global System for Mobile Communications. *See* GSM
 GPS, enabling, 106
 gray box testing, 289
 grey hats, 158–159
 Group Policy, using for configuration management, 360–361
 GSM, 104
 GSM Association (GSMA), 104

Guardians of Peace (GOP), 161
Guide for Applying the Risk Management Framework to Federal Information Systems, 394–395
Guide for Conducting Risk Assessments, 244–245
Guide to Integrating Forensic Techniques, 443
Guide to Integrating Forensic Techniques into Incident Response, 447
Guide to Intrusion Detection and Prevention Systems (IDPS), 279
Guide to IPsec VPNs, 134
Guide to Protecting the Confidentiality of Personally Identifiable Information (PII), 387
Guidelines on PDA Forensics, 447
Guidelines on Security and Privacy in Public Cloud Computing, 145

H

hackers, 158
 hardening systems, 163, 314–315
 hardware
 asset management, 390–391
 losing, 162
 security policies, 411
 hardware tokens, 34
 hashing, 10–12, 135, 183, 467, 469
 algorithms, 469–472
 tools, 472–473
 verifying a hash, 472–473
 heuristic-based detection, 226–228
 See also viruses
 HIDS. *See* host-based intrusion detection systems (HIDS)
 High Speed Packet Access (HSPA), 105
 HIPAA, 148, 354, 388–389, 452
 HMAC, 471
 hoaxes, 218–219
 honeynets, 286
 honeypots, 285–286
 host-based intrusion detection systems (HIDS), 278–279
 hot sites, 13, 425–426
 Hotspot questions, 5, 7
 HTTP, 86
 HTTPS, 9–10, 86
 hybrid detection methods, 283
 HyperText Transfer Protocol. *See* HTTP
 HyperText Transfer Protocol Secure (HTTPS). *See* HTTPS
 hypervisor, 140–141
 improperly configured, 143

I

ICMP, 83, 290
 IDEA, 478
 identification, 27–28
 reviewing, 38
 identity management, 55–56
 account lockout policies, 57
 de-provisioning, 58
 entitlement, 57–58
 maintenance and entitlement, 57
 principle of least privilege, 56, 57–58
 provisioning and authorization, 56–57

- identity proofing, 56
 - Identity Theft Resource Center (ITRC), 452, 453
 - IDSs. *See* intrusion detection systems (IDSs)
 - IEEE 802.11, 97
 - See also* wireless technologies
 - IEEE 802.3, 73
 - IGMP, 83–84
 - IMAP4, 87–88
 - impact, 249–250
 - impersonation, 195
 - inappropriate usage, 263
 - incident response, 262–263, 439–440
 - containment, eradication, and recovery, 266
 - detection and analysis, 265–266
 - discovering an incident, 440
 - escalating an incident, 440–441
 - first responders and preserving the scene, 442–443
 - implementing countermeasures, 442
 - implementing feedback loops, 441–442
 - post-incident activity, 267
 - preparation, 263–265
 - reporting an incident, 441
 - responding to an incident, 440
 - and security policies, 410
 - indirect losses, vs. direct losses, 413
 - information rights management, 381–382
 - information security programs, goals of, 9
 - Information Technology Infrastructure Library. *See* ITIL
 - Information Technology Laboratory. *See* ITL
 - Infrastructure-as-a-Service (IaaS), 144, 145
 - injection attacks, 178–180
 - input validation
 - client side vs. server side, 175
 - as a countermeasure, 174–175
 - insider attacks, 160–161
 - vs. APTs, 161
 - Institute of Electrical and Electronics Engineers (IEEE), 10
 - integrated services digital network. *See* ISDN
 - integrity, 10–12, 465, 485
 - enforcing with hashing, 469–474
 - See also* cryptography
 - internal testing, 288
 - International Information Systems Security Certification Consortium, Inc. *See* (ISC)²
 - International Organization for Standardization (ISO), 67
 - Internet, 91
 - communications, 120–121
 - connections in rural areas, 122
 - Internet Control Message Protocol. *See* ICMP
 - Internet Group Message Protocol. *See* IGMP
 - Internet Message Access Protocol version 4. *See* IMAP4
 - Internet Protocol security. *See* IPsec
 - Internet Small Computer System Interface. *See* iSCSI
 - internetwork trust architectures, 91–97
 - intranets, 91
 - intrusion detection systems (IDSs), 263, 265, 275–276
 - alerts, 276–277
 - and defense in depth, 279
 - firewall logs, 352
 - host-based intrusion detection systems (HIDS), 278–279
 - network-based IDS (NIDS), 276, 277–278
 - intrusion prevention systems (IPSs), 263, 279–281
 - detection methods, 282–283
 - iOS, 323–324
 - IP addresses, 125
 - public and private, 93–94
 - IPsec, 88–89, 132, 133–134, 492
 - IPSs. *See* intrusion prevention systems (IPSs)
 - IPv4
 - vs. IPv6, 79–80
 - link local addresses, 94
 - IPv6
 - vs. IPv4, 79–80
 - link local addresses, 94
 - irises, 36
 - IRM. *See* information rights management
 - ISACA, 356
 - (ISC)², 1
 - Associate designation, 4
 - chapters, 4
 - Code of Ethics, 2–3
 - iSCSI, 141
 - ISDN, 121
 - ISDN DSL (IDSL), 121
 - ISO. *See* International Organization for Standardization (ISO)
 - ITIL, 429
 - ITL, 429
- ## J
- jailbreaking, 324
 - job rotation, 451–452
 - JPEG, 72
- ## K
- KeePass, 31, 32
 - Kerberos, 39–40
 - See also* single sign-on authentication
 - Kernel Patch Protection, 216
 - Kernel PatchGuard, 216
 - Key Areas of Knowledge, 5
 - key composition, 501
 - Key Distribution Centers (KDCs), 39
 - key escrow, 499–500
 - key rotation, 501
 - keyloggers, 215
 - keys, composing and rotating, 475–476
 - known-plaintext attacks, 502
 - KryptoKnight, 42
- ## L
- L2F, 132
 - L2TP, 132, 133
 - labeling data, 374
 - LAND, 165
 - See also* denial of service (DoS) attacks
 - LANs, 73, 93, 115
 - comparing switches and routers, 115–117
 - network segmentation, 117–119
 - secure device management, 120
 - See also* VLANs

LaPadula, Leonard J., 53–54
 Layer 2 Forwarding. *See* L2F
 Layer 2 Tunneling Protocol. *See* L2TP
 least privilege, 13–14, 232–233
 link local addresses, 94
 Linksys, 102
 LizaMoon SQL injection attack, 179, 213
 load-balancing clusters, 330
 local area network denial attack. *See* LAND
 local area networks. *See* LANs
 logic bombs, 215
 logical access controls, 47
 Logical Link Control sublayer, 70
 logs. *See* audit logging
 loss, defined, 243
 losses, 413
 LTE, 105

M

MAC address filtering, 43, 103
 MAC addresses, 116–117
 macro viruses, 210
 See also viruses
 Madoff, Bernie, 450
 maillog, 351
 Makwana, Rajendrasinh, 160
 malicious code, 207
 backdoors and trapdoors, 217–218
 countermeasures, 223–234
 keyloggers, 215
 logic bombs, 215
 malware hoaxes, 218–219
 mobile code, 216
 ransomware, 213–214
 Regin, 219–221
 remote access tools (RATs), 218
 rootkits, 215–216
 scareware, 211–212, 213
 spyware, 218
 Trojan horses, 211
 viruses, 207–210
 worms, 210
 malware, 263, 322
 anti-analysis malware, 443
 blocking, 128
 delivering via e-mail, 223
 delivering via malvertising, 222
 delivering via USB drives, 223
 drive-by downloads, 221–222
 forwarding, 162
 hoaxes, 218–219
 jailbreaking, 324
 newsletters and bulletins about, 227
 See also GameOver Zeus
 management/administrative security controls, 311–312
 Mandatory Access Control (MAC) model, 52–55
 mandatory vacation policies, 450–451
 Mandiant, 159
 Mandiant U.S.A. Cyber Security Ransomware, 213
 man-in-the-middle attacks, 171–172
 MANs, 93
 marking data, 374
 maximum acceptable outage, 420
 maximum tolerable downtime (MTD), 420
 maximum tolerable outage (MTO), 420
 MCI WorldCom, 450
 MD5. *See* Message Digest 5
 Md5sum.exe, 473
 Media Access Control (MAC) filtering, 43, 103
 Media Access Control (MAC) sublayer, 69, 70
 media access units (MAUs), 78
 mesh topology, 78–79
 Message Digest 5, 35, 135, 470–471
 metamorphic viruses, 209
 See also viruses
 metropolitan area networks. *See* MANs
 Microsoft
 Cybercrime Center, 167
 Digital Crimes Unit (DCU), 167
 and virtualization, 140
Minimum Security Requirements for Federal Information and Information Systems, 390
 MITM. *See* man-in-the-middle attacks
 MITRE Corporation, 235
 mobile code, 216
 mobile device management (MDM), 322
 mobile devices, protecting, 106
 mobile sites, 427
 Moneypak Virus, 213
 monitoring systems
 analyzing results, 283–284
 detecting unauthorized changes, 284–286
 detection systems and logs, 284
 See also intrusion detection systems (IDSs); intrusion prevention systems (IPSs); wireless intrusion detection systems; wireless intrusion prevention systems
 Morris worms, 210
 MPEG, 72
 MS-CHAPv1 and MS-CHAPv2, 135
 multifactor authentication, 37–38
 multipartite viruses, 209
 See also viruses
 multistation access units (MSAUs), 78
 Mydoom, 217

N

NAC, 138–140
 Nash, Michael J., 55
 NAT, 94–95, 98, 123
 and proxy server logs, 351–352
 National Cyber Security Division, 235, 430
 National Institute of Standards and Technology.
 See NIST
 National Vulnerability Database (NVD), 254
 NDR, 82
 near field communication. *See* NFC
 Nessus, 290
 Netgear, 102
 network access control. *See* NAC
 network access, in security policies, 410

Network Address Translation. *See* NAT
 Network Discovery Protocol. *See* NDP
 Network File System. *See* NFS
 network interface cards (NICs), 69
 network segmentation, 117–119
 network topologies, 73–79
 network-based IDS (NIDS), 276, 277–278
 New Technology File System (NTFS), 47
 newsletters, 197
 next-generation firewalls, 128
 NFC, 105–106
 NFS, 47, 87, 141
 NIDS. *See* network-based IDS (NIDS)
 NIST, 244, 428–429
 Nmap, 290
 no operation (NOOP) commands, 177
 nonce, 135, 173
 non-Discretionary Access Control (non-DAC) models, 49
 nonrepudiation, 16–17, 484–485
 NTFS, extended attributes, 220

O

objects, 44, 47
 offline authentication, 43
 OLAP, 385–386
 OLTP, 385–386
 on-demand scanning, 224
 One-Time Password (HOTP) protocol, 34–35
 One-Time Password in Everything (OPIE), 35
 one-time passwords, 29, 34–35
 one-way trust, 96
 online analytical processing. *See* OLAP
 online transaction processing. *See* OLTP
 Open Shortest Path First. *See* OSPF
 Open Systems Interconnection model. *See* OSI model
 Openfiler, 141
 operating systems
 keeping up to date, 231
 logs, 348–349
 operational security controls, 312–313
 organization mission statement, 409
 OSI model
 Application layer, 72
 compared to TCP/IP model, 72–73
 Data Link layer, 69–70
 Network layer, 70
 overview, 67–68
 Physical layer, 68–69
 Presentation layer, 72
 Session layer, 71–72
 Transport layer, 70–71
 OSPF, 87

P

packet sniffers. *See* protocol analyzers
 packet-filtering firewalls, 125–127
 padded cells, 286
 palms, 35
 PANs, 93, 104
 PAR, 134
 partial knowledge testing. *See* gray box testing
 passphrases, 30
 password attacks, 182–185
 Password Authentication Protocol. *See* PAP
 password protection, 106
 Password-Based Key Derivation Function 2 (PBKDF2), 478
 passwords
 asynchronous dynamic passwords, 34
 auditing, 355
 audits, 31
 cognitive, 29–30, 56
 dynamic, 29, 34
 guidelines, 30–31
 one-time, 29, 34–35
 Password Policy settings, 31–32
 salting, 473–474
 and security policies, 411
 static, 29
 strong, 30, 33
 synchronous dynamic passwords, 34
 wireless device administrator password, 101
 PAT, 95
 patching systems, 163, 318
 applying patches, 319
 auditing systems, 319
 documenting patches, 320
 evaluating patches, 318
 testing patches, 318–319
 PBX, 122–123
 PCI DSS, 148, 355, 468
 control objectives, 357
 requirements, 356–358
 and security policies, 413
 PEAR, 135
 Pearson VUE test centers, 1
 creating an account, 2
 penetration tests, 294
 permission creep, 50
 persistent data, 445
 personal area networks. *See* PANs
 Personal Identity Verification (PIV) cards, 38
 personally identifiable information. *See* PII
 pharming, 195–196
 PHI, 15, 388–389
 regulatory requirements, 468
 phishing, 185–187
 and APT1, 288
 and drive-by downloads, 187
 responding to, 162
 spear phishing and whaling, 187–188
 See also smishing; vishing
 phreaks, 162
 physical access controls, auditing, 358
 physical security controls, 313
 See also security policies
 physical security operations, 58–59
 piggybacking, 194
 PII, 15, 387–388, 452–453
 regulatory requirements, 468
 ZIP codes, 455

- ping of death, 165
 - See also* denial of service (DoS) attacks
 - ping sweeps, 170
 - PINs, 33
 - PKI. *See* public key infrastructure (PKI)
 - plaintext, 466
 - Platform-as-a-Service (PaaS), 144, 145
 - Point-to-Point Tunneling Protocol. *See* PPTP
 - Police Virus, 213
 - polymorphic viruses, 209
 - See also* viruses
 - POODLE attacks, 86, 501
 - Port Address Translation. *See* PAT
 - port mirroring, 120
 - port scans, 170, 277
 - ports, 125
 - comparing ports and protocol numbers, 91
 - mapping well-known ports to protocols, 89–90
 - PPTP, 132, 133
 - presheared key. *See* PSK
 - pretexting, 193
 - Pretty Good Privacy (PGP), 490
 - preventive controls, 307–308
 - principle of least privilege, 56, 57–58
 - privacy, 15
 - California Online Privacy Protection Act (OPPA), 456–457
 - Children’s Online Privacy Protection Act (COPPA), 456
 - in cloud computing, 147
 - E-Privacy Directive, 454
 - European directives, 454–455
 - legal issues, 452–453
 - Safe Harbor program, 454
 - security fundamentals, 15
 - Social Security numbers, 455–456
 - ZIP codes, 455
 - private branch exchange. *See* PBX
 - Protected EAP. *See* PEAP
 - protected health information. *See* PHI
 - protocol analyzers, 167–170
 - protocols, 79–89, 125
 - common protocol numbers, 91
 - common protocols and their port numbers, 90
 - mapping well-known ports to, 89–90
 - See also individual protocols*
 - provisioning, 56–57
 - proximity cards, 35
 - proxy servers, 123–125, 346
 - logs, 351–352
 - PSK, 101
 - PSTN, 120–121
 - public key infrastructure (PKI), 492
 - alternative certificate trusts, 500–501
 - certificate authorities (CAs), 495–496
 - certificate trust chain, 497, 500–501
 - certificates, 492–495
 - key escrow, 499–500
 - revoking certificates, 498
 - Trusted Root Certification Authorities, 497–498
 - validating certificates, 498–499
 - public switched telephone network. *See* PSTN
 - purging media, 378
- ## Q
- qualitative analysis, 256–258
 - quantitative analysis, 254–256
- ## R
- radio frequency identification. *See* RFID
 - radio frequency interference (RFI), 69
 - RADIUS, 100, 136–137
 - RAID, 141, 325–329
 - See also* fault tolerance
 - rainbow tables, 183–184
 - ransomware, 213–214
 - RC4, 478, 479
 - real-time scanning, 224
 - recovery agents, 499–500
 - recovery controls, 310
 - recovery point objective, 421
 - recovery time objective (RTO), 420
 - redundant connections, 13, 330–331
 - redundant disks, 12
 - redundant servers, 12
 - redundant sites, 13
 - Regin, 219–221
 - regulatory requirements
 - cryptography, 468
 - HIPAA, 388–389
 - personally identifiable information (PII), 387–388
 - Sarbanes-Oxley (SOX) Act of 2002, 389
 - training, 390
 - remote access solutions
 - authentication, 134
 - CHAP, 135
 - Diameter, 137
 - EAP, 135–136
 - IPsec, 133–134
 - L2F, 132
 - L2TP, 133
 - MS-CHAPv1 and MS-CHAPv2, 135
 - overview, 130–131
 - PAP, 134
 - PPTP, 133
 - RADIUS, 136–137
 - risks and vulnerabilities, 131
 - SSH, 132
 - TACACS+, 138
 - TLS, 134
 - traffic shaping, 138
 - tunneling protocols, 131–132
 - remote access tools (RATs), 218
 - Remote Authentication Dial-in User Service. *See* RADIUS
 - Remote SPAN (RSPAN), 120
 - remote wipe, 106
 - replay attacks, 173
 - response plans, 317
 - restoration planning, 424
 - See also* business continuity plans (BCPs)

- retention requirements, 377
 - retinas, 36
 - RFC 1918, 93
 - RFC 4193, 94
 - RFID, 105, 391
 - RIPv2, 87
 - risk, defining, 243–245
 - risk assessments, 254
 - address findings, 262
 - Guide for Conducting Risk Assessments*, 244–245
 - qualitative analysis, 256–258
 - quantitative analysis, 254–256
 - steps, 258–262
 - risk management, 250–251
 - certification and accreditation, 394–395
 - identifying assets, 252
 - residual risk, 251–252
 - risk register, 253–254
 - risk visibility and reporting, 253
 - in security policies, 410
 - risks and vulnerabilities, total risk, 244
 - rogue access points, 191
 - rogueware, 211–212
 - Role-based Access Control (Role-BAC) model, 49–50
 - Rombertik, 443
 - rootkits, 215–216
 - ROT13, 475
 - routers, 115–117
 - Cisco routers, 120
 - routing, 98
 - Routing Information Protocol version 2. *See* RIPv2
 - routing protocols, 87
 - RSA, 481
 - Rule-based Access Control (Rule-BAC) model, 51
 - Rustok botnet, 167
 - See also* botnets
- ## S
- Safe Harbor program, 454
 - safeguards. *See* security controls
 - salami attacks, 171
 - salting, 122, 473–474
 - SAML, 41–42
 - See also* single sign-on authentication
 - sandboxing, 232, 325
 - as a countermeasure, 176–177
 - sanitizing media, 378
 - SANS Institute, 430
 - SANS Internet Storm Center, 227, 430
 - Sarbanes-Oxley (SOX) Act, 389
 - satellite connections, 121
 - scanners, 231
 - scareware, 211–212, 213
 - SCCM. *See* ConfigMgr
 - scheduled scanning, 224
 - script kiddies, 161–162
 - SDN and NFV Market Size Report*, 143
 - Secure European System for Applications in a Multivendor Environment. *See* SESAME
 - Secure Real-Time Transport Protocol. *See* SRTP
 - Secure Shell. *See* SSH
 - Secure Sockets Layer. *See* SSL
 - SecureAuth Identity Provider (IdP), 44
 - Security and Privacy Controls for Federal Information Systems and Organizations*, 306, 310
 - Security Assertion Markup Language. *See* SAML
 - security audits, 354–355
 - auditing passwords, 355
 - auditing physical access controls, 358
 - PCI DSS requirements, 356–358
 - security awareness, 469
 - Security Considerations in the System Development Life Cycle*, 395
 - security controls
 - change control and configuration management, 317–318
 - classes of controls, 311–313
 - combining control goals and classes, 313–314
 - compensating controls, 309–310
 - control families, 312
 - corrective controls, 309
 - cost-benefit analyses, 304–305
 - detective controls, 308
 - deterrent controls, 310
 - directive controls, 310
 - goals of, 307–310
 - hardening systems, 314–315
 - life cycle, 305–306
 - management/administrative, 311–312
 - operational, 312–313
 - overview, 303–304
 - physical, 313
 - policies, standards, procedures, and guidelines, 315–317
 - preventive controls, 307–308
 - recovery controls, 310
 - response plans, 317
 - technical, 312
 - testing patches, fixes, and updates, 318–320
 - See also* device security
 - security fundamentals
 - AAAs of security, 17–18
 - accountability, 18
 - defense in depth, 15–16
 - due care, 19–20
 - due diligence, 19
 - least privilege, 13–14
 - nonrepudiation, 16–17
 - privacy, 15
 - separation of duties, 14–15
 - Security Guide for Interconnecting Information Technology Systems*, 392
 - security identifiers (SIDs), 48
 - security information and event management. *See* SIEM
 - security operations
 - asset management, 390–391
 - classifying data, 371–374
 - regulatory requirements, 387–390
 - system development life cycle (SDLC), 395–398
 - See also* certification and accreditation; data management policies

- security organizations, 428–430
- security policies, 316
 - auditing, 355
 - becoming more common, 413–414
 - characteristics, 408
 - easy-to-read language, 415
 - enforcing, 412
 - flyers and posters, 416
 - overview, 407–408
 - policy awareness, 415
 - stages, 408
 - training sessions, 416
 - updating, 416–417
 - value of, 412–413
 - warning banners, 415–416
- segmentation, 117–119
- Sender Policy Framework (SPF) records, 229
- separation of duties, 14–15
- Server Message Blocks. *See* SMBs
- service level agreements (SLAs), 147–148
- service set identifiers (SSIDs), 102–103
- SESAME, 42
- session hijacking, 172–173
- SHA-1, 471
- sha1sum.exe, 472–473
- SHA-2, 471
- SHA-3, 471
- shielded twisted pair (STP) cable, 69
- shortened links, 231–232
- shoulder surfing, 195
- SIEM, 286–287
- signature-based detection, 225–226, 282
 - keeping AV signatures up to date, 228
 - See also* intrusion prevention systems (IPSs); viruses
- Simple Integrity Axiom, 54
- Simple Mail Transfer Protocol. *See* SMTP
- Simple Network Management Protocol. *See* SNMP
- single loss expectancy (SLE), 254–255
- single sign-on authentication, 38
 - federated access, 40–41
 - Kerberos, 39–40
 - KryptoKnight, 42
 - SAML, 41–42
 - SESAME, 42
- smart cards, 33, 38
- SMBs, 131–132
- S/MIME, 484–490
- smishing, 188
- SMTP, 87–88
- smurf attacks, 173–174
- sniffers. *See* protocol analyzers
- sniffing attacks, 167–170
- SNMP, 84
- social engineering, 183
 - dumpster diving, 195
 - impersonation, 195
 - overview, 193–194
 - pharming, 195–196
 - pretexting, 193
 - shoulder surfing, 195
 - social networking attacks, 196
 - tailgating, 194
 - user awareness as a countermeasure, 196–197
- social network usage, 381
- social networking attacks, 196
- Social Security numbers, 455–456
- software security, 233–234
 - as a countermeasure, 174
- software tokens, 34–35
- Software-as-a-Service (SaaS), 144–145
- software-defined networking (SDN), 142, 143
- something you are, 35–37
- something you have, 33–35
- something you know, 29–33
- Sony, 388
- SOX. *See* Sarbanes-Oxley (SOX) Act
- spam, 185
 - filtering, 128
 - filters, 229
- spear phishing, 187–188
- spoofing, 163
- spyware, 218
- SQL, 385
- SQL injection attacks, 178–179, 213
- SRTT, 122
- SSCP certification
 - Candidate Information Bulletin (CIB), 5
 - Common Body of Knowledge (CBK), 3, 5
 - maintaining, 7–8
 - requirements, 1
 - years of experience, 3–4
- SSCP certification exam
 - fee, 2
 - future question types, 5–7
 - Key Areas of Knowledge, 5
 - passing the exam, 4–5
 - registering for, 1–2
- SSH, 85–86, 132, 484
- SSL, 86–87, 483
- SSO. *See* single sign-on authentication
- star Integrity Axiom, 54
- star topology, 76–77
- stateful inspection firewalls, 127
- statement of accountability, 409
- static passwords, 29
- stealth viruses, 209
 - See also* viruses
- steganography, 490–491
- storage media, data management policies, 376–377
- stored procedures, 179
- stream ciphers, 476
- strong passwords, 30, 33
- structural threat sources, 246
- Structured Query Language. *See* SQL
- subjects, 44, 45–46
- subnet addresses, 125
- sulog, 350
- Switched Port Analyzer (SPAN), 120

switches, 115–117
 symmetric encryption, 467
 3DES, 478
 Advanced Encryption Standard, 477
 Bcrypt, 478
 Blowfish, 478
 comparing block and stream ciphers, 476
 composing and rotating keys, 475–476
 DES, 478
 IDEA, 478
 overview, 474
 Password-Based Key Derivation Function 2 (PBKDF2), 478
 RC4, 478, 479
 ROT13, 475
 SYN flood attacks, 164
 See also denial of service (DoS) attacks
 synchronous dynamic passwords, 34
 syslog, 350
 System Center Configuration Manager (ConfigMgr), 319
 system development life cycle (SDLC), 395–398
 Systems Security Certified Practitioner certification.
 See SSCP certification

T

tabletop exercises, 424–425
 TACACS+, 138
 tailgating, 194
 TCP/IP model, compared to OSI model, 72–73
 technical security controls, 312
 telecommunications
 converged communications, 123
 Internet communications, 120–121
 Internet connections in rural areas, 122
 overview, 120
 securing phones, 122–123
 VoIP, 122
 Telnet, 85
 Temporal Key Integrity Protocol. *See* TKIP
 Terminal Access Controller Access Control System+.
 See TACACS+
 testing and drills, 424–425
 TFTP, 84–85
 third-party connections. *See* extranets
 threat events, 245
 identifying, 246–248
 threat sources, 245
 identifying, 245–246
 threats, defined, 243
 thumbprints, 35
 Time-based One-Time Password (TOTP)
 protocol, 34
 TKIP, 100
 TLS, 86–87, 132, 134, 481–483
 token ring topology, 77–78
 topography, vs. topology, 74
 topologies, 73–79
 vs. topography, 74

traffic shaping, 138
 training, 197, 325, 390
 security awareness, 469
 sessions, 416
 transitive trust, 96–97
 Transmission Control Protocol (TCP), 70–71
 Transport Layer Security. *See* TLS
 trapdoors, 217–218
 tree topology, 77
 Trivial File Transfer Protocol (TFTP), 71
 Trivial FTP. *See* TFTP
 Trojan horses, 211
 Trojan Reveton, 213
 trust chain, 497, 500–501
 trust relationships, 96–97
 Trusted Computer System Evaluation Criteria (TCSEC), 394
 Trusted Platform Module (TPM), 324
 Trusted Root Certification Authorities, 497–498
 tunneling protocols, 88, 131–132
 twisted pair cable, 69
 two-way trust, 96

U

unauthorized connections, 285
 unauthorized data access, 162, 263
 unified threat management (UTM) devices, 277
 United States Computer Emergency Readiness Team.
 See US-CERT
 United States Cyber Command (USCYBERCOM), 160
 United States Government Configuration Baseline (USGCB), 359–360
 unpatched systems, and zero day exploits, 189
 unrecoverable error rate (URE), 329
 unshielded twisted pair (UTP) cable, 69
 unsolicited commercial e-mail (UCE), 185
 URL filtering, 128, 351
 USB devices
 delivering malware, 223
 security controls, 322
 US-CERT, 429–430
 user awareness, 325
 as a countermeasure, 196–197
 increasing, 163
 User Datagram Protocol (UDP), 70, 71, 173–174
 users, 375

V

virtual appliances, 141
 infected, 143
 virtual local area networks. *See* VLANs
 virtual machines (VMs), 140, 323
 virtual networks, improperly configured, 144
 virtual private networks. *See* VPNs
 virtualization
 attacks and countermeasures, 143–144
 continuity and resilience, 142
 and device security, 323

virtualization (*Cont.*)
 overview, 140
 separation of data plane and control plane, 142
 shared storage, 141
 software-defined networking (SDN), 143
 terminology, 140–141

Virus Bulletin, 227

viruses, 207–210
See also antivirus software; heuristic-based detection;
 signature-based detection

vishing, 188

VLANs, 117–119
See also LANs

VM escape, 144, 323

VMware, 140

Voice over Internet Protocol. *See* VoIP

VoIP, 122

volatile data, 445

VPNs, 130–138

vulnerabilities, 245, 248–249
 defined, 243

vulnerability assessments, 287–288
 attempting to exploit vulnerabilities, 295
 discovery, 295
 false positives, 293–294
 fingerprinting, 295
 identifying vulnerabilities, 295
 infrastructure security configurations, 291
 order of steps, 291–293
 penetration tests, 294
 performing, 295
 reporting results to management, 296
 results analysis, 293–294
 test types, 288–289
 vulnerability-scanning tools, 289–291
 written permission, 293, 295

W

WANs, 93
 traffic shaping, 138

WAP, 97

wardriving, 101, 192

warm sites, 13, 427

warning banners, 415–416

web spiders, 185

whaling, 187–188

white box testing, 288

white hats, 158–159

whitelists, 125, 234, 283

wide area networks. *See* WANs

WIDS. *See* wireless intrusion detection systems

Wi-Fi Alliance, 10, 99

Wi-Fi Protected Access (WPA), 10, 99, 100, 101
 WPA cracking attacks, 192–193

Wi-Fi Protected Access 2 (WPA2), 99, 100, 101

Wilson, David, 54–55

WiMAX, 105

Win32/Zbot. *See* GameOver Zeus

Windows Local Group Policy Editor, 31–32

Windows Management Instrumentation (WMI), 391

Windows Server Update Services (WSUS), 319

WIPS. *See* wireless intrusion prevention systems

Wired Equivalent Privacy (WEP), 10, 99, 101

wireless attacks and countermeasures, 190–191

wireless device administrator password, 101

wireless intrusion detection systems, 191, 283

wireless intrusion prevention systems, 283

wireless service set identifiers, 102–103

wireless standards, 97

wireless technologies, 97–106

wireless transmissions, 69

WireLurker, 324

Wireshark, 99, 168–169

Worldwide Interoperability for Microwave Access.
See WiMAX

worms, 210

WPA cracking attacks, 192–193

WPA-Enterprise, 100–101

WPA-Personal, 100

WPA2, as a countermeasure, 190–191

WPA2-Enterprise, 100–101

WPA2-Personal, 100

X

XACML, 52

XSRF. *See* cross-site request forgery

XSS. *See* cross-site scripting

Y

years of experience, 3–4, 5

Z

zero day exploits, 188–190, 222

zero knowledge testing. *See* black box testing

Zeus. *See* GameOver Zeus

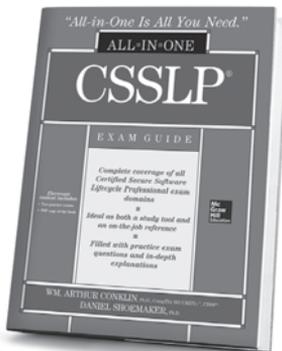
ZIP codes, as PII, 455

zombies, 165–167

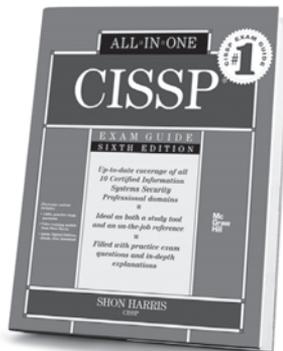
Complete coverage of today's top

IT SECURITY

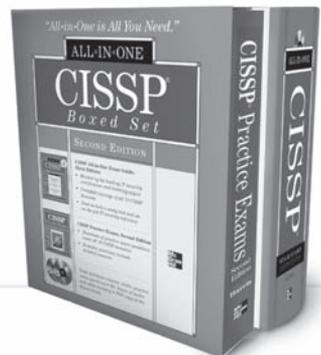
certification exams



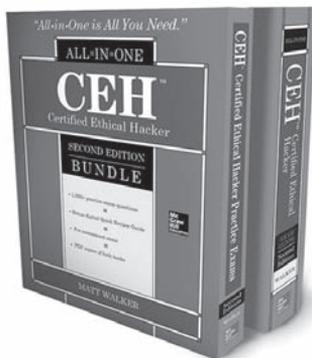
0-07-176026-1 • \$60.00



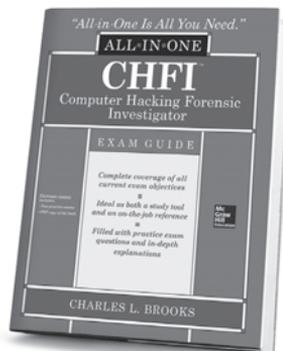
0-07-178174-9 • \$80.00



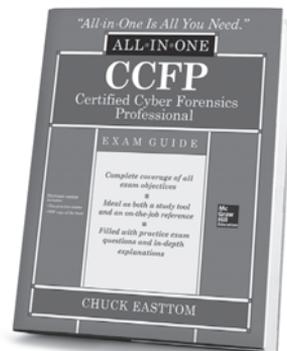
0-07-179308-9 • \$105.00



0-07-183557-1 • \$70.00



0-07-183156-8 • \$50.00



0-07-183976-3 • \$60.00

LICENSE AGREEMENT

THIS PRODUCT (THE "PRODUCT") CONTAINS PROPRIETARY SOFTWARE, DATA AND INFORMATION (INCLUDING DOCUMENTATION) OWNED BY MCGRAW-HILL EDUCATION AND ITS LICENSORS. YOUR RIGHT TO USE THE PRODUCT IS GOVERNED BY THE TERMS AND CONDITIONS OF THIS AGREEMENT.

LICENSE: Throughout this License Agreement, "you" shall mean either the individual or the entity whose agent opens this package. You are granted a non-exclusive and non-transferable license to use the Product subject to the following terms:

- (i) If you have licensed a single user version of the Product, the Product may only be used on a single computer (i.e., a single CPU). If you licensed and paid the fee applicable to a local area network or wide area network version of the Product, you are subject to the terms of the following subparagraph (ii).
- (ii) If you have licensed a local area network version, you may use the Product on unlimited workstations located in one single building selected by you that is served by such local area network. If you have licensed a wide area network version, you may use the Product on unlimited workstations located in multiple buildings on the same site selected by you that is served by such wide area network; provided, however, that any building will not be considered located in the same site if it is more than five (5) miles away from any building included in such site. In addition, you may only use a local area or wide area network version of the Product on one single server. If you wish to use the Product on more than one server, you must obtain written authorization from McGraw-Hill Education and pay additional fees.
- (iii) You may make one copy of the Product for back-up purposes only and you must maintain an accurate record as to the location of the back-up at all times.

COPYRIGHT; RESTRICTIONS ON USE AND TRANSFER: All rights (including copyright) in and to the Product are owned by McGraw-Hill Education and its licensors. You are the owner of the enclosed disc on which the Product is recorded. You may not use, copy, decompile, disassemble, reverse engineer, modify, reproduce, create derivative works, transmit, distribute, sublicense, store in a database or retrieval system of any kind, rent or transfer the Product, or any portion thereof, in any form or by any means (including electronically or otherwise) except as expressly provided for in this License Agreement. You must reproduce the copyright notices, trademark notices, legends and logos of McGraw-Hill Education and its licensors that appear on the Product on the back-up copy of the Product which you are permitted to make hereunder. All rights in the Product not expressly granted herein are reserved by McGraw-Hill Education and its licensors.

TERM: This License Agreement is effective until terminated. It will terminate if you fail to comply with any term or condition of this License Agreement. Upon termination, you are obligated to return to McGraw-Hill Education the Product together with all copies thereof and to purge all copies of the Product included in any and all servers and computer facilities.

DISCLAIMER OF WARRANTY: THE PRODUCT AND THE BACK-UP COPY ARE LICENSED "AS IS." MCGRAW-HILL EDUCATION, ITS LICENSORS AND THE AUTHORS MAKE NO WARRANTIES, EXPRESS OR IMPLIED, AS TO THE RESULTS TO BE OBTAINED BY ANY PERSON OR ENTITY FROM USE OF THE PRODUCT, ANY INFORMATION OR DATA INCLUDED THEREIN AND/OR ANY TECHNICAL SUPPORT SERVICES PROVIDED HEREUNDER, IF ANY ("TECHNICAL SUPPORT SERVICES"). MCGRAW-HILL EDUCATION, ITS LICENSORS AND THE AUTHORS MAKE NO EXPRESS OR IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE WITH RESPECT TO THE PRODUCT. MCGRAW-HILL EDUCATION, ITS LICENSORS, AND THE AUTHORS MAKE NO GUARANTEE THAT YOU WILL PASS ANY CERTIFICATION EXAM WHATSOEVER BY USING THIS PRODUCT. NEITHER MCGRAW-HILL EDUCATION, ANY OF ITS LICENSORS NOR THE AUTHORS WARRANT THAT THE FUNCTIONS CONTAINED IN THE PRODUCT WILL MEET YOUR REQUIREMENTS OR THAT THE OPERATION OF THE PRODUCT WILL BE UNINTERRUPTED OR ERROR FREE. YOU ASSUME THE ENTIRE RISK WITH RESPECT TO THE QUALITY AND PERFORMANCE OF THE PRODUCT.

LIMITED WARRANTY FOR DISC: To the original licensee only, McGraw-Hill Education warrants that the enclosed disc on which the Product is recorded is free from defects in materials and workmanship under normal use and service for a period of ninety (90) days from the date of purchase. In the event of a defect in the disc covered by the foregoing warranty, McGraw-Hill Education will replace the disc.

LIMITATION OF LIABILITY: NEITHER MCGRAW-HILL EDUCATION, ITS LICENSORS NOR THE AUTHORS SHALL BE LIABLE FOR ANY INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES, SUCH AS BUT NOT LIMITED TO, LOSS OF ANTICIPATED PROFITS OR BENEFITS, RESULTING FROM THE USE OR INABILITY TO USE THE PRODUCT EVEN IF ANY OF THEM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL APPLY TO ANY CLAIM OR CAUSE WHATSOEVER WHETHER SUCH CLAIM OR CAUSE ARISES IN CONTRACT, TORT, OR OTHERWISE. Some states do not allow the exclusion or limitation of indirect, special or consequential damages, so the above limitation may not apply to you.

U.S. GOVERNMENT RESTRICTED RIGHTS: Any software included in the Product is provided with restricted rights subject to subparagraphs (c), (1) and (2) of the Commercial Computer Software-Restricted Rights clause at 48 C.F.R. 52.227-19. The terms of this Agreement applicable to the use of the data in the Product are those under which the data are generally made available to the general public by McGraw-Hill Education. Except as provided herein, no reproduction, use, or disclosure rights are granted with respect to the data included in the Product and no right to modify or create derivative works from any such data is hereby granted.

GENERAL: This License Agreement constitutes the entire agreement between the parties relating to the Product. The terms of any Purchase Order shall have no effect on the terms of this License Agreement. Failure of McGraw-Hill Education to insist at any time on strict compliance with this License Agreement shall not constitute a waiver of any rights under this License Agreement. This License Agreement shall be construed and governed in accordance with the laws of the State of New York. If any provision of this License Agreement is held to be contrary to law, that provision will be enforced to the maximum extent permissible and the remaining provisions will remain in full force and effect.