



CISSP Certification



Roberta Bragg, CISSP

Training Guide

CISSP TRAINING GUIDE

Copyright© 2003 by Que Publishing

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

International Standard Book Number: 0-7897-2801-x

Library of Congress Catalog Card Number: 2002110896

Printed in the United States of America

First Printing: November 2002

04 03 02 01 4 3 2 1

Trademarks

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Que Publishing cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The author and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the CD or programs accompanying it.

PUBLISHER

Paul Boger

EXECUTIVE EDITOR

Jeff Riley

ACQUISITIONS EDITOR

Jeff Riley

DEVELOPMENT EDITOR

Ginny Bess Munroe

MANAGING EDITOR

Thomas Hayes

PROJECT EDITOR

Tonya Simpson

PRODUCTION EDITORS

Megan Wade

Michael Dietsch

INDEXER

John Sleeva

PROOFREADER

Juli Cook

TECHNICAL EDITORS

Guy Bruneau

Lawrence S. Paccone

Patrick Ramseier

TEAM COORDINATOR

Rosemary Lewis

MULTIMEDIA DEVELOPER

Michael Hunter

INTERIOR DESIGNER

Louisa Kluczniak

COVER DESIGNER

Anne Jones

PAGE LAYOUT

Cheryl Lynch

GRAPHICS

Tammy Graham

Oliver Jackson

Contents at a Glance

<i>Introduction</i>	1
---------------------------	---

PART I Exam Preparation

1 Access Control Systems and Methodology	13
2 Telecommunications and Network Security	61
3 Security Management and Practices	175
4 Applications and Systems Development Security	235
5 Cryptography	307
6 Security Architecture and Models	335
7 Operations Security	381
8 Business Continuity Planning and Disaster Recovery Planning	439
9 Law, Investigation, and Ethics	493
10 Physical Security	529

PART II Final Review

Fast Facts	575
Study and Exam Prep Tips	621
Practice Exam	625

PART III Appendixes

A Glossary	667
B Overview of the Certification Process	687
C What's on the CD-ROM	689
D Using the <i>PrepLogic Practice Tests, Preview Edition</i> Software	691
<i>Index</i>	697







Table of Contents

PART I: Exam Preparation

1 Access Control Systems and Methodology	13
Introduction	17
Accountability	18
Access Control Techniques	19
Discretionary Access Control	20
Mandatory Access Control	21
Lattice-Based Access Control	22
Rule-Based Access Control	25
Role-Based Access Control	26
Access Control Lists	27
Access Control Administration	27
Account Administration	28
Access Control Models	29
Bell-LaPadula	30
Biba	32
Summary of BLP and Biba	33
Liptner's Lattice	33
Non-Inference Models	33
Identification and Authentication Techniques	34
Passwords	35
One-Time Passwords	35
Challenge Response	36
Biometrics	36
Tickets	36
Single Sign-On	37
Access Control Methodologies	37
Centralized/Remote Authentication Access Controls	38
Decentralized Access Control	38

VI CISSP TRAINING GUIDE

Methods of Attacks	40
Brute-Force	41
Denial-of-Service	42
Spoofing	42
Sniffing	43
Monitoring	43
Intrusion Detection	43
Intrusion Prevention	46
How Intrusion Detection Works	46
Penetration Testing	48
Penetration Testing Versus Security Assessments	49
Ethical Issues	49
Performing a Penetration Test	50
Common Tools	51
Exercises	55
Review Questions	55
Exam Questions	56
Answers to Review Questions	57
Answers to Exam Questions	59
2 Telecommunications and Network Security	61
Introduction	67
The Open Systems Interconnection Model	68
The OSI Layers	70
OSI Summary	77
Network Characteristics and Topologies	78
Coax	79
Unshielded Twisted Pair	82
Fiber Optic	84
Wireless	87
Network Topologies	89
Linear Bus Topology	89
Star Topology	91
Ring Topology	92
Tree Topology	93
Mesh Topology	93
LAN and WAN Technologies	94
Ethernet	95

Token-Ring and FDDI	98
Attached Resource Computer Network	99
LAN Devices	99
Hubs and Repeaters	99
Switches and Bridges	100
VLANs	101
Routers	103
Firewalls	104
Gateways and Proxies	110
WAN Technologies	110
Dedicated Connections	111
Circuit-Switched Connections	113
Packet-Switched Connections	113
Cell-Switched Connections	114
WAN Services	114
WAN Devices	118
Providing Remote Access Capabilities	119
Client-Based Dial-in Remote Access	119
Using Tunneling As a Security Method	120
Virtual Private Networks	121
Remote Access Authentication	124
Networking Protocols	125
Transmission Control Protocol/Internet Protocol	125
Reviewing TCP and UDP	129
Protecting the Integrity, Availability, and Confidentiality of Network Data	130
The CIA Triad	130
Security Boundaries and Translating Security Policy to Controls	132
Trusted Network Interpretation	133
Network Layer Security Protocols	135
Transport Layer Security Protocols	136
Application Layer Security Protocols	136
Network Monitoring and Packet Sniffers	137
Intrusion Detection	139
Intrusion Response	141
Network Address Translation	142
Transparency	144
Hash Totals	145
Email Security	146

VIII CISSP TRAINING GUIDE

Facsimile and Printer Security	147
Common Attacks and Countermeasures	147
Fault Tolerance and Data Restoration	155
Managing Network Single Points of Failure	158
Topology Failures	159
Exercises	165
Review Questions	166
Exam Questions	167
Answers to Review Questions	169
Answers to Exam Questions	171
3 Security Management and Practices	175
Introduction	179
Defining Security Principles	180
CIA: Information Security's Fundamental Principles	180
Privacy	183
Identification and Authentication	184
Nonrepudiation	188
Accountability and Auditing	188
Documentation	190
Security Management Planning	191
Risk Management and Analysis	192
Risk Analysis	194
Identifying Threats and Vulnerabilities	195
Asset Valuation	196
Qualitative Risk Analysis	202
Countermeasure Selection and Evaluation	203
Tying It Together	204
Policies, Standards, Guidelines, and Procedures	205
Information Security Policies	206
Setting Standards	209
Creating Baselines	210
Guidelines	210
Setting and Implementing Procedures	210
Examining Roles and Responsibility	212
Management Responsibility	213
User Information Security Responsibilities	213
IT Roles and Responsibilities	214
Other Roles and Responsibilities	214

Understanding Protection Mechanisms	215
Layering	216
Abstraction	217
Data Hiding	217
Encryption	217
Classifying Data	218
Commercial Classification	219
Government Classification	220
Criteria	221
Creating Procedures for Classifying Data	221
Employment Policies and Practices	222
Background Checks and Security Clearances	222
Employment Agreements, Hiring, and Termination	223
Job Descriptions	225
Job Rotation	225
Managing Change Control	226
Hardware Change Control	226
Software Change Control	227
Security Awareness Training	227
Exercises	230
Review Questions	230
Exam Questions	230
Answers to Review Questions	232
Answers to Exam Questions	232
4 Applications and Systems Development Security	235
Introduction	239
Software Applications and Issues	240
Challenges of Distributed and Nondistributed Environments	241
Database and Data Warehousing Issues	249
Storage and Storage Systems	256
Knowledge-Based Systems	261
Web Services and Other Examples of Edge Computing	262
Attacking Software	266
Attacks Against Password Databases	266
Denial-of-Service and Distributed Denial-of-Service Attacks	267
Spoofing	269
Miscellaneous Attacks	270

X CISSP TRAINING GUIDE

Illegitimate Use of Legitimate Software	272
Network Software	273
Understanding Malicious Code	274
So, Who's a Hacker? What's Malicious Code?	275
What Protection Does Antivirus Software Provide?	277
Implementing System Development Controls	277
System Development Lifecycle	278
Security Control Architecture	283
Best Practices	285
Using Coding Practices That Reduce System Vulnerability	286
Software Development Methodologies	286
Impacting Security Through Good Software Design and Coding Practices	292
Exercises	300
Review Questions	301
Exam Questions	301
Answers to Review Questions	302
Answers to Exam Questions	303
5 Cryptography	307
Introduction	310
Uses of Cryptography	310
Confidentiality	310
Integrity	311
Authentication	311
Nonrepudiation	312
Cryptographic Concepts, Methodologies, and Practices	313
Symmetric Algorithms	313
Asymmetric Algorithms	315
Message Authentication	316
Hash Functions	316
Digital Signatures	317
Key Length	317
One-Time Ciphers	318
PKI and Key Management	318
Methods of Attack	319
General Attacks	320
Specific Attacks	322

Exercises	329
Review Questions	329
Exam Questions	329
Answers to Review Questions	330
Answers to Exam Questions	331
6 Security Architecture and Models	335
Introduction	338
Requirements for Security Architecture and Models	340
Security Models	342
Bell-LaPadula	342
Biba	345
Clark-Wilson Model	346
Access Control Lists	347
A Review of the Security Models	347
Security System Architecture	348
Reference Monitor	348
Open Versus Closed Systems	350
Security Principles	351
Security Modes	352
Labels Versus Access Control Lists	353
Covert Channel	354
Information System Security Standards	355
TCSEC—The Orange Book and the Rainbow Series	356
Information Technology Security Evaluation Criteria	360
Common Criteria	362
What Is Common Criteria?	363
A Comparison of the Orange Book, ITSEC, and Common Criteria	370
IPSec	370
Uses for IPSec	371
Architectural Components of IPSec	372
Exercises	377
Review Questions	377
Exam Questions	377
Answers to Review Questions	378
Answers to Exam Questions	379

XII CISSP TRAINING GUIDE

7 Operations Security	381
Introduction	385
Examining the Key Roles of Operations Security	387
Identify Resources to Be Protected	387
Identifying Privileges to Be Restricted	388
Identifying Available Controls and Their Types	389
Control Types	391
Describing the OPSEC Process	391
The Roles of Auditing and Monitoring	395
Using Logs to Audit Activity and Detect Intrusion	396
Detecting Intrusions	399
Penetration Testing Techniques	403
Developing Countermeasures to Threats	408
Risk Analysis	408
Threats	409
Countermeasures	411
Establishing Countermeasures for Employee-Related Threats	412
Including Countermeasures in Hiring and Firing/Exit Practices	414
Gruntling Program	415
Countermeasures for Common Internet-Based Threats	416
Countermeasures to Physical Threats	417
The Role of Administrative Management	418
Concepts and Best Practices	420
Privileged Operation Functions	421
Understanding Antiviral Controls	423
Protecting Sensitive Information and Media	425
Change Management Control	427
Exercises	432
Answer to Exercise	433
Review Questions	434
Exam Questions	434
Answers to Review Questions	435
Answers to Exam Questions	436
8 Business Continuity Planning and Disaster Recovery Planning	439
Introduction	444
What Are the Disasters That Interrupt Business Operation?	445
Quantifying the Difference Between DRP and BCP	448

Examining the Business Continuity Planning Process	450
Determining the Plan's Scope	451
Business Impact Assessment	452
The BIA Process	458
Developing Operational Plans	458
Implementing the Plan	464
Testing the Plan	464
Maintaining the Plan	465
Defining Disaster Recovery Planning	466
Recovering Data Processing	467
Restoring Data Processing	472
Developing a Backup Strategy	472
Backup Procedures and Policy	474
Vital Records Program	477
Hardware Backups	478
Alternative Sites	478
Exercises	485
Review Questions	485
Exam Questions	486
Answers to Review Questions	487
Answers to Exam Questions	488
9 Law, Investigation, and Ethics	493
Introduction	496
Fundamentals of Law	497
Intellectual Property Law	498
Sale and Licensing	499
Privacy Law	500
Government Regulations	502
Criminal Law and Computer Crime	503
Computer Security Incidents	505
Advance Planning	506
Computer Crime Investigation	507
Legal Evidence	509
Credibility or Weight of Evidence	510
Proof of Authenticity	511
Hearsay	511
Best Evidence Rule	511

XIV CISSP TRAINING GUIDE

Chain of Evidence	512
The Fourth Amendment	513
Computer Forensics	513
Computer Ethics	517
Exercises	523
Review Questions	523
Exam Questions	523
Answers to Review Questions	525
Answers to Exam Questions	526
10 Physical Security	529
Introduction	532
Classifying Assets to Simplify Physical Security Discussions	533
Vulnerabilities	535
Selecting, Designing, Constructing, and Maintaining a Secure Site	538
Site Location and Construction	539
Physical Access Controls	540
Power	544
Environmental Controls: Air Conditioning, Humidity, and Temperature	547
Water Exposure Problems	548
Fire Prevention and Protection	549
Tape and Media Library Retention Policies	553
Document (Hard-Copy) Libraries	555
Waste Disposal	556
Physical Intrusion Detection	559
Exercises	565
Review Questions	567
Exam Questions	568
Answers to Review Questions	569
Answers to Exam Questions	570

PART II: Final Review

Fast Facts	575
Domain 1, "Access Control"	576
Accountability	576
Access Controls	576

Access Control Administration	576
Access Control Models	576
Identification and Authentication Techniques	577
Remote Authentication Access Control	577
Centralized Versus Decentralized Access Control	577
Methods of Attack	577
Monitoring	578
Penetration Testing and System Assessment	578
Domain 2, "Network Security and Telecommunications"	578
ISO/OSI Seven-Layer Model	578
Network Cabling	579
Network Topologies	580
LAN and WAN Technologies	580
Network Devices	581
Firewalls	581
Gateways and Proxies	582
Connection Speeds and Types	582
Connections	582
WAN Services	583
WAN Devices	583
Remote Access	583
TCP/IP	584
Common Network Attacks and Countermeasures	584
Fault Tolerance	584
Domain 3, "Security Management and Practices"	585
CIA Triad	585
Privacy	585
Identification and Authentication	585
Auditing	586
Risk Management and Analysis	586
Qualitative Risk Analysis	586
Cost-Effectiveness of a Countermeasure	587
Responses to Risk Analysis	587
Policies	587
Protection Mechanisms	587
Data Classification	587
Government Data Classification	588
Domain 4, "Applications and Systems Development Security"	588
Centralized and Distributed Systems	588
Risks for Centralized Computer Systems	589

XVI CISSP TRAINING GUIDE

Database Management Systems	589
SANs	591
Web Services Issues	591
More on Attacks	591
Malicious Code	592
System Development Models	592
Security Control Architecture	593
Software Development Methodologies	594
Coding for Security	594
Domain 5, “Cryptography”	594
Uses of Cryptography	594
Cryptographic Methods and Algorithms	595
Attacks Against Encryption	596
Domain 6, “Security Architecture and Models”	597
Examining the Differences Between Government and Industry Models	597
Security Models	597
Security Architecture	597
Covert Channels	599
Information Security Standards	599
Uses for IPSec	603
Domain 7, “Operations Security”	603
Roles of Operations Security	603
Role of Auditing Monitoring	605
Intrusion Detection	605
Penetration Testing Techniques	605
Countermeasures to Threats	605
Employee Job Duties and Risks to Systems	606
The Role of Administrative Management	608
Principles of OPSEC	608
Antiviral Controls	609
Management of Sensitive Data	609
Change Management Control	609
Domain 8, “Business Continuity Planning and Disaster Recovery Planning”	610
Mandated Plans	610
Differences Between DRP and BCP	610
Business Continuity Planning Process	610
Business Impact Assessment	610
Operations Plan	610
Insurance	610
Testing the Plan	611

CONTENTS XVII

Maintenance	611
Disaster Recovery Planning	611
Antidisaster Procedures	612
Backup Issues to Consider	612
Alternative Sites	612
Domain 9, “Law, Investigation, and Ethics”	613
Criminal, Civil, and Administrative Law	613
Intellectual Property Law	613
Sales and Licensing	613
Privacy	614
Federal Laws	614
Criminal Law	615
Computer Crime Investigation	615
Evidence	616
Fourth Amendment to the U.S. Constitution	616
Forensics	616
Ethics	617
Domain 10, “Physical Security”	617
Classification of Assets	617
Countermeasure to Theft	618
Site Location and Construction	618
Physical Access Controls	618
Power	618
Environmental Controls	619
Water Exposure Problems	619
Fire Prevention	619
Fire Extinguishers	619
Tape and Media Retention Policy	620
Waste Disposal	620
Study and Exam Prep Tips	621
Learning As a Process	621
Study Tips	622
Study Strategies	622
Pretesting Yourself	623
Exam Prep Tips	623
Putting It All Together	623
Practice Exam	625
Exam Questions	626
Answers to Exam Questions	654

PART III: Appendixes

A Glossary	667
B Overview of the Certification Process	687
Description of the Path to Certification	687
About the Certification Program	687
C What's on the CD-ROM	689
<i>PrepLogic Practice Tests, Preview Edition</i>	689
Exclusive Electronic Version of Text	689
D Using the <i>PrepLogic Practice Tests, Preview Edition</i> Software	691
Exam Simulation	691
Question Quality	691
Interface Design	691
Effective Learning Environment	691
Software Requirements	692
Installing <i>PrepLogic Practice Tests, Preview Edition</i>	692
Removing <i>PrepLogic Practice Tests, Preview Edition</i> from Your Computer	692
Using <i>PrepLogic Practice Tests, Preview Edition</i>	692
Starting a Practice Test Mode Session	693
Starting a Flash Review Mode Session	693
Standard <i>PrepLogic Practice Tests, Preview Edition</i> Options	693
Time Remaining	694
Your Examination Score Report	694
Review Your Exam	694
Get More Exams	694
Contacting PrepLogic	695
Customer Service	695
Product Suggestions and Comments	695
License Agreement	695
Index	697

Preface

PAPER CHASE

From *Security Watch*, Monday, June 3, 2002

By Roberta Bragg

I'll admit it. I hold a paper cert. I attended bartending school and passed the exam, but I'm not employed as a bartender, nor do I have any real experience. Would any of you believe that some piece of paper qualifies me to take over at your local watering hole?

How is this different from being a paper MCSE? The market, that's how. After attending bartending school, you're qualified to apply to places so desperate for help they'll take someone with no experience. The meaning is clear when you list that credential.

Being a paper cert holder in some other professional circles isn't a good thing. The proliferation of paper MCSEs—those who have passed the exams but have no experience—has hurt us all. (Don't get me wrong; a paper MCSE can earn respect in my book if she goes out and gets the experience and doesn't claim that the paper alone makes her an experienced professional.) And now we're faced with something even more dangerous: paper CISSPs. CISSP stands for Certified Information System Security Professional, a vendor-neutral security cert widely heralded as *the* security certification for serious security professionals. The certification is managed by (ISC)². To obtain the title, you must not only pass a grueling examination, but also prove you have four years of information security experience and sign a statement of ethics. (More information is available at www.isc2.org and my article at <http://certcities.com/editorial/exams/print.asp?EditorialsID=25>.)

The requirement for experience before obtaining a certification isn't unusual. A physician must have an internship before she can become an MD, teachers have their student teaching, and CPAs must be employed in the field. A CISSP candidate must have four years of experience. Starting this year, each applicant must have a supervisor's signature to back up his or her claims. Those with inadequate experience will not be allowed to take the exam—and yes, some statements will be audited. Because of this required proof of experience, I believe the certification will continue to mean something. Unfortunately, some schools, and a great number of you, apparently do not.

If four years experience is required, how can there be paper CISSPs? Recently, I've received reports that training schools are offering CISSP bootcamps and encouraging participants to lie about their real-world experience when they apply for CISSP certification.

The existence and use of bootcamps is a controversial subject. Those against it say no one can learn the information in that short a period and retain it. Others say experienced technical people should use them to polish their knowledge and pass the exams. As far as I'm concerned, bootcamps may or may not be a good thing. I don't like them when they create paper-certified people, but I see their usefulness for those with experience.

My problem lies with those training centers that would

- ◆ Seek to cheat students out of money by guaranteeing the cert to those not meeting the CISSP experience requirement
- ◆ Cheat legitimate certification holders by cheapening the certification

- ◆ Cheat the industry by presenting poorly prepared people as applicants for jobs in which the security of their networks lie
- ◆ Cheat all participants by encouraging students to lie on the experience form and get someone to validate the lie with a signature

Liars will lie, and cheaters will cheat; we'll never find a way to catch all of them, and maybe I'm a little naive to think I can change that. I guess it's up to you then. You, the ethical IT folks, will have to get involved.

Bootcamp companies should stop overpromising, fully inform prospective students about the CISSP requirements, and stop encouraging them to lie. Hiring managers should investigate experience claims. And if you're considering taking the low road to your security certification, stop. Do it right, and get the experience first.

Enough is enough. It's time we in the IT profession get a grip on our ethical behavior and let others know we'll hold them to the same standards.

*Reprinted with permission from Security Watch
(<http://lists.101com.com/NLS/pages/main.asp>), ©
2002 101 Communications, LLC.*

Roberta Bragg, MCSE, CISSP, runs her company Have Computer Will Travel, Inc., out of a notebook carrying case. She's an independent consultant specializing in security, operating systems, and databases. Send her your questions or comments at <mailto:roberta.bragg@mcpmag.com>.

CISSP BOOT CAMPS: CAVEAT EMPTOR

From *Security Watch*, Monday, August 12, 2002

By Marc Thompson

We share Roberta Bragg's concern regarding the proliferation of Certified Information System Security Professional (CISSP) "boot camps," as she wrote in the June 3 issue of *Security Watch*.

First, it should be noted that these boot camps have no affiliation with the International Information Systems Security Certification Consortium (ISC)², the non-profit organization of security executives that manages the CISSP credential. It also should be noted that (ISC)² provides the only officially sanctioned training for CISSP candidates through its education arm, (ISC)² Institute.

In addition, (ISC)² Institute occasionally offers official review seminars via approved training partners. These partners host the program, while (ISC)² Institute provides all the instructors and coursework. Affiliated training programs are listed on the (ISC)² Web site.

That said, the consortium has no legal right to stop someone from claiming to train for CISSP certification. It's the same with other IT training programs from Cisco or Microsoft where third parties can offer to prepare candidates for a credential.

All CISSP candidates should tread carefully when dealing with these boot camps. They should be forewarned that if they're found to be lying regarding their past work experience, as Roberta claims some boot camps are encouraging, they'll lose their certification for violating the CISSP "Code of Ethics" they're required to sign prior to taking the exam and are legally committed to.

In the past year, (ISC)² has taken several steps to minimize the ability of candidates to misrepresent their work experience, including random audits of applications and requiring a candidate to obtain an endorsement of their professional experience by a CISSP. The endorser attests that the candidate's assertions regarding professional experience are true to the best of their knowledge, and that the candidate is in good standing within the information security industry.

In addition, if a CISSP candidate attends a boot camp that utilizes materials from the actual test, as some boot camps claim, the candidate will also be in violation of the Code of Ethics and will lose their CISSP certification. Also, candidates shouldn't believe that a boot camp can increase CISSP exam pass rates, as several claim. As a matter of policy, (ISC)² has never published its pass rates, so there is no way for a boot camp to legitimately claim high pass rates.

The key difference between the boot camps and (ISC)² Institute training is fundamental: The institute's goal is to provide an extensive overview of the Common Body of Knowledge (CBK), the compendium of information security practices and standards compiled and continually updated by (ISC)² and used as the basis for the CISSP exam.

(ISC)² Institute's instructors, who are all CISSPs trained by the consortium, allow CISSP candidates to understand the level of their knowledge of the CBK's 10 domains for later study before taking the CISSP exam. Some of the institute's instructors have been training CISSP candidates for 5 years or more, while many boot camps have often only been operating for a few months. Again, it's the only training (ISC)² recommends.

We want to reassure Roberta and other concerned CISSPs that (ISC)² is making every effort to ensure that the certification remains the "gold standard" in the information security industry. We fully support her call for ethical behavior among all IT professionals.

Marc Thompson is Vice President of the (ISC)² Institute (<http://www.isc2.org/>).

Reprinted with permission from Security Watch (<http://lists.101com.com/NLS/pages/main.asp>), © 2002 101 Communications, LLC.

About the Authors

Roberta Bragg, CISSP, MCSE, and the original Security Evangelist, is a veteran of more than 25 years in IT. Her technical experience ranges from programming to systems administration and Windows network security design. She is an internationally acclaimed author and lecturer on Windows security.

Scott Barman is currently an information security and systems architecture analyst for The MITRE Corporation (www.mitre.org) working with the MITRE team to help the IRS modernize its IT infrastructure. He has been involved with information security for almost 20 years, nurturing the evolution of systems and their security requirements for commercial organizations and government agencies. Since the explosion of the Internet, and prior to joining MITRE, he has focused on various areas of security and policy development for many organizations in the Washington, D.C. area. Scott earned his undergraduate degree from the University of Georgia and a Master of Information Systems Management with a concentration in information security management from Carnegie Mellon University (www.mism.cmu.edu).

Philip Fites has worked for more than 34 years in informatics, from computer operations to business and project management. His current focus includes information systems security theory and practice. Since the early 1980s, a lifelong interest in information security has been transformed into a commitment to research on integrity and other issues of security in information systems, combined with a practical focus on applying his expertise to help clients clarify and achieve security objectives.

Philip holds a bachelor of science in mathematics and an M.B.A. and studied for a Ph.D. in computing science at Queen's University. He is coauthor of *Control and Security of Computer Information Systems*, *The Computer Virus Crisis*, and *Information Systems Security: A Practitioner's Reference*, and he has published a number of works on various topics in computer security, software research, and educational planning methodology in various professional and industry publications. He has served as a director and president of the International Information Systems Security Certification Consortium (ISC)². He is a member of the Standards Council of Canada's Canadian Advisory Committee on Information Technology.

Wesley J. Noonan is currently a senior quality assurance representative with BMC Software, Inc. (www.bmc.com) working on its network management product line. Wes got his start in the United States Marine Corps working on its Banyan VINES network and has spent the past 10 years building, maintaining, and securing corporate networks ranging in size from 25 to 25,000 users. Wes is also an active trainer, developing and teaching his own custom, Cisco-based routing and switching curriculum. His certifications include MCSE, CCNA, CCDA, and NNCSS.

Benjamin Wright, recognized the world over as one of the leading lawyers in e-commerce, is the founding author of *The Law of Electronic Commerce*, a comprehensive book on the legality of electronic transactions, published by Aspen Law & Business. A graduate of Georgetown University Law Center, he is an independent attorney practicing computer security and e-commercial law in Dallas, Texas.

Since 1988, he has delivered more than 500 speeches on e-commerce, privacy, and computer security and has been quoted in publications around the globe, from the *Wall Street Journal* to the *Sydney Morning Herald*. On May 26, 2001, he was featured in the 30-minute documentary *The Cutting Edge Technology Report: Electronic Signatures*, nationally broadcast on CNBC.

About the Technical Reviewers


Guy Bruneau, GSEC, GCIA, GCUX is a senior security consultant with InfoPeople Security Solutions, Inc. He works within InfoPeople's security practice assisting clients with their managed security services, computer intrusion detection operations and deployment, network security auditing, incident response and reporting, and so on. He has firsthand knowledge in the use and hardening of Cisco Secure IDS, Shadow IDS, and Snort IDS, among others.

He has been a SANS instructor and speaker and is the author of the IDIC course *Introduction to Logfile Analysis*. He is an authorized SANS Unix security grader and is presently serving as the chair of the SANS GIAC Certified Intrusion Analyst Advisory Board. He is the author of the OS hardened Shadow IDS platform based on NSWC's Shadow version 1.7 (available at <http://www.whitehats.ca>). In his spare time, he has worked as a technical reviewer for New Riders Publishing.

Lawrence S. Paccone is a principal national/systems security analyst at Northrop Grumman Information Technology TASC. As both a technical lead and project manager, he has worked in the Internet and network/systems security arena for more than 8 years.

He has been the technical lead for several network security projects supporting a government network/systems security research and development laboratory. Prior to that, he worked for 5 years at The Analytical Sciences Corporation (TASC) as a national security analyst assessing conventional military force structures. He has an M.S. in information systems, an M.A. in international relations, and a B.A. in political science. He has completed eight professional certifications in network and systems security, internetworking, wide area networking, Cisco routing/switching, Unix, and Windows NT. He also has been a technical editor for eight IT security books that are currently in publication.

Patrick "Swissman" Ramseier, CCNA, CISSP is a systems engineer at OKENA, makers of the StormSystem Intrusion Prevention System. OKENA has been delivering breakthrough security software products that proactively preserve the operational integrity of applications and host systems. OKENA StormSystem is a system of seamlessly integrated security products that act in unison to prevent existing and unknown attacks without relying on attack signatures. Patrick started out as a Unix system administrator. Over the past 14 years, he has been involved with corporate-level security design; architecture reviews; vulnerability assessments; VPN support; physical, network, and operating system security (Unix-Solaris, Linux, BSD, and Windows NT/2000); training; research; and post- and pre-sales. He has a B.A. in business and is working concurrently on his master's and doctorate in computer science.






Dedication

This one's for the postman.

Acknowledgments

Thanks go to the eds, who left me alone on this one.



We Want to Hear from You!

As the reader of this book, *you* are our most important critic and commentator. We value your opinion and want to know what we're doing right, what we could do better, what areas you'd like to see us publish in, and any other words of wisdom you're willing to pass our way.

As a publisher for Que, I welcome your comments. You can email or write me directly to let me know what you did or didn't like about this book--as well as what we can do to make our books better.





Please note that I cannot help you with technical problems related to the *topic* of this book. We do have a User Services group, however, where I will forward specific technical questions related to the book.





When you write, please be sure to include this book's title and author as well as your name, email address, and phone number. I will carefully review your comments and share them with the author and editors who worked on the book.

Email: feedback@quepublishing.com

Mail: Jeff Riley
Que Publishing
201 West 103rd Street
Indianapolis, IN 46290 USA

For more information about this book or another Que title, visit our Web site at www.quepublishing.com. Type the ISBN (excluding hyphens) or the title of a book in the Search field to find the page you're looking for.





How to Use This Book

Que Certification has made an effort in its Training Guide series to make the information as accessible as possible for the purposes of learning the certification material. Here, you have an opportunity to view the many instructional features that have been incorporated into the books to achieve that goal.

CHAPTER OPENER

Each chapter begins with a set of features designed to allow you to maximize study time for that material.

List of Objectives: Each chapter begins with a list of the objectives as stated by the exam's vendor.

Objective Explanations: Immediately following each objective is an explanation of it, providing context that defines it more meaningfully in relation to the exam. Because vendors can sometimes be vague in their objectives list, the objective explanations are designed to clarify any vagueness by relying on the authors' test-taking experience.

OBJECTIVES

Understand the principles of security management.


- ▶ In understanding information security management, there are a number of principles you need to know to create a managed security program. These principles go beyond firewalls, encryptions, and access control. They are concerned with the various aspects of managing the organization's information assets in areas such as privacy, confidentiality, integrity, accountability, and the basics of the mechanisms used in their management.

Know what management's responsibility is in the information security environment.

- ▶ Management cannot just decree that the systems and networks will be secure. They must take an active role in setting and supporting the information security environment. Without management support, the users will not take information security seriously.

Understand risk management and how to use risk analysis to make information security management decisions.

- ▶ Managing security is the management of risk. Knowing how to assess and manage risk is key to an information security management program.



3

CHAPTER

Security Management
and Practices

OUTLINE

Introduction	179	Identify from Whom It Is Being Protected	209
		Setting Standards	209
		Creating Baselines	210
		Guidelines	210
		Setting and Implementing Procedures	210
Defining Security Principles	180	Examining Roles and Responsibility	212
CIA: Information Security's Fundamental Principles	180	Management Responsibility	213
Confidentiality	181	User Information Security Responsibilities	213
Integrity	182	IT Roles and Responsibilities	214
Availability	183	Other Roles and Responsibilities	214
Privacy	183	Understanding Protection Mechanisms	215
Identification and Authentication	184	Layering	216
Passwords	185	Abstraction	217
Nonrepudiation	188	Data Hiding	217
Accountability and Auditing	188	Encryption	217
Keystroke Monitoring	189	Classifying Data	218
Protecting Audit Data	190	Commercial Classification	219
Documentation	190	Government Classification	220
		Criteria	221
Security Management Planning	191	Creating Procedures for Classifying Data	221
Risk Management and Analysis	192		
Risk Analysis	194		
Identifying Threats and Vulnerabilities	195		
Asset Valuation	196		
Qualitative Risk Analysis	202		
Countermeasure Selection and Evaluation	203		
Tying It Together	204		

Chapter Outline: Learning always gets a boost when you can see both the forest and the trees. To give you a visual image of how the topics in a chapter fit together, you will find a chapter outline at the beginning of each chapter. You will also be able to use this for easy reference when looking for a particular topic.

STUDY STRATEGIES

- ▶ Operations security covers a lot of ground. From management of equipment to management of people, the topics it relates to have no end. One of the challenges of understanding this broad topic is its reliance on the underlying technology. Concepts and best practices will not make sense if you do not understand how computers, networks, programs, data centers, and businesses work. If you do not already have some experience with them, try to find someone who does and ask them to help you understand. Spending some time with Chapter 2, "Telecommunications and Network Security," will help as well.
- ▶ If your background is not in technology, mastering the material in the domains of Security Management, Telecommunications and Networking Security, Disaster Recovery and Business Continuity, and Application and Systems Development is essential. Study them first, and this chapter will be easier. If your background is technical, do not be frustrated by the light treatment of technical content here. Operations security is more concerned with the big picture than with the intimated details of how to configure or code.
- ▶ Whatever your current job description, whatever your background, use your knowledge of this domain to do two things:
 - First, see what you can determine about operations security at your organization. If you aren't working in a directly related area, don't be surprised if you can't find out much. Good operations security is transparent; that is, it reveals little about the specifics of its activities.
 - Second, operations security principles can be applied to things other than computer operations. Military organizations have long used these principles to improve their prospects of success. Practice these principles on your activities on the Internet. What can someone learn about you while you're online? How might that be used to their advantage? What can you do to diminish the amount of information that can be gleaned from your activities?

Study Strategies: Each topic presents its own learning challenge. To support you through this, Que Certification has included strategies for how to best approach studying in order to retain the material in the chapter, particularly as it is addressed on the exam.

INSTRUCTIONAL FEATURES WITHIN THE CHAPTER

These books include a large amount and different kinds of information. The many different elements are designed to help you identify information by its purpose and importance to the exam and also to provide you with varied ways to learn the material. You will be able to determine how much attention to devote to certain elements, depending on what your goals are. By becoming familiar with the different presentations of information, you will know what information will be important to you as a test-taker and which information will be important to you as a practitioner.

Objective Coverage Text: In the text before an exam objective is specifically addressed, you will notice the objective is listed to help call your attention to that particular material.

Note: Notes appear in the margins and contain various kinds of useful information, such as tips on the technology or administrative practices, historical background on terms and technologies, or side commentary on industry issues.

NOTE

Effectiveness and Functionality of Countermeasures Choosing a countermeasure for the amount of cost is a pure business way of analyzing risk. However, as security professionals, we understand that regardless of the cost, the countermeasure is not worth using unless it protects the asset. Information security professionals should work with business people to select the most effective countermeasure that will function to properly protect the asset.

As technology grew, the information assets became less centralized and management had the problem of maintaining the integrity of the network and the information being used on the systems on the networks. Although there is a move to try to centralize management of servers and information security, information security management needs to take into account everywhere the information assets touch.

Network computing has brought new paradigms to the sharing of information. Using technologies such as virtual private networks (VPNs) and extranets, organizations can forge new types of relationships based on sharing information assets. These partnerships have organizations connecting their networks to share information in a way that was unheard of as recently as 10 years ago. Managers planning these partnerships also should keep in mind how to maintain the security of other information assets not involved in those agreements. Both organizations should consider undergoing a risk analysis specific to the connectivity required for this partnership to provide appropriate protections.

RISK MANAGEMENT AND ANALYSIS

Understand risk management and how to use risk analysis to make information security management decisions.

Risk management is the process of assessing risk and applying mechanisms to reduce, mitigate, or manage risks to the information assets. Risk management is not about creating a totally secure environment. Its purpose is to identify where risks exist, the probability that the risks could occur, the damage that could be caused, and the costs of securing the environment. Even if there is a risk to information assets, risk management can determine that it would cost more to secure the asset than if it was damaged or disclosed.

Risk management is not as straightforward as finding the risk and quantifying the cost of loss. Because risks can come from varying sources, an information asset can have several risks. For example, sales data stored on a network disk has the risk of

- ◆ Unauthorized access from internal or external users
- ◆ Loss from a software or hardware failure
- ◆ Inaccessibility because of a network failure

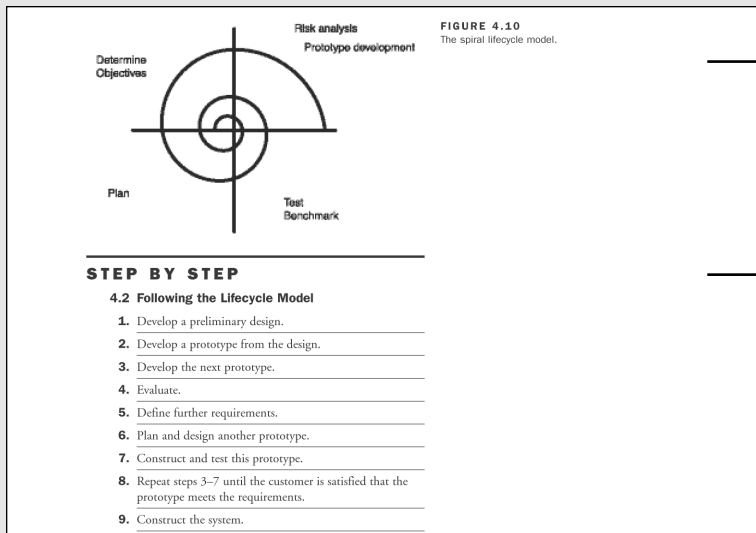


Figure: To improve readability, the figures have been placed in the margins wherever possible so they do not interrupt the main flow of text.

Step by Step: Step by Steps are hands-on tutorial instructions that walk you through a particular task or function relevant to the exam objectives.

NOTE **Effectiveness and Functionality of Countermeasures** Choosing a countermeasure for the amount of cost is a pure business way of analyzing risk. However, as security professionals, we understand that regardless of the cost, the countermeasure is not worth using unless it protects the asset. Information security professionals should work with business people to select the most effective countermeasure that will function to properly protect the asset.

IN THE FIELD

PHYSICAL ATTACK PARAMETERS

Several observations regarding site selection and building construction have been made in this chapter. In situations of national security or where terrorism might be a factor, careful attention must be given to measures that will lessen vulnerability to physical attacks.

Many sources (for example, Van Nostrand Reinhold's *Computer Security Risk Management* and RCMP *Security Information Publications # 3*) provide information about typical times for various methods of physical penetration and using various tools.

REVIEW BREAK

Tying It Together

Risk assessment tells the organization what the risks are; it is up to the organization to determine how to manage the risks. Risk management is the trade-off an organization makes regarding that risk. You should remember that not every risk could be mitigated. It is the job of management to decide how that risk is handled. In basic terms, the choices are

- ▶ **Do nothing**—If you do this, you must accept the risk and the potential loss if the threat occurs.
- ▶ **Reduce the risk**—You do this by implementing a countermeasure and accepting the residual risk.

NOTE **Residual Risk** This is the value of the risk after implementing the countermeasure.

In the Field Sidebar: These more extensive discussions cover material that perhaps is not as directly relevant to the exam, but which is useful as reference material or in everyday practice. In the Field may also provide useful background or contextual information necessary for understanding the larger topic under consideration.

Review Break: Crucial information is summarized at various points in the book in lists or tables. At the end of a particularly long section, you might come across a Review Break that is there just to wrap up one long objective and reinforce the key points before you shift your focus to the next section.

CASE STUDIES

Case Studies are presented throughout the book to provide you with another, more conceptual opportunity to apply the knowledge you are developing. They also reflect the “real-world” experiences of the authors in ways that prepare you not only for the exam but for application in your job. In each Case Study, you will find similar elements: a description of a Scenario, the Essence of the Case, and an extended Analysis section.

Analysis: This is a lengthy description of the best way to handle the problems listed in the Essence of the Case. In this section, you might find a table summarizing the solutions, a worded example, or both.

CASE STUDY: DOES BUSINESS CONTINUITY WORK?

ESSENCE OF THE CASE

- ▶ A business continuity plan was in place; however, the unique way in which employees responded to a disaster ensured this company's continuation and subsequent successes.

SCENARIO

Yes (and the better your plans, the more likely it is). In the wake of the 9/11 attack on the World Trade Center, many businesses did not survive. But many did. The World Trade Center offices of bond trading giant Cantor Fitzgerald LP were destroyed, and 180 of its 733 employees were killed. However, Cantor was ready to trade two days later—in time for the September 13 reopening of U.S. Treasury markets.

According to an article in the December 13, 2001 issue of *Computerworld* (<http://www.cnn.com/2001/TECH/industry/12/13/redundancy.rebound.idg/index.html>) and information on the company's Web site (www.espeed.com), Cantor was able to do so because of built-in redundancy provided by its business-to-business online marketplace and IT services group, eSpeed (www.espeed.com), and because of the efforts of remaining eSpeed employees based in the U.S. and London. eSpeed had duplicated its IT services in a similar data center in the U.S. and was working toward

Essence of the Case: A bulleted list of the key problems or issues that need to be addressed in the Scenario.

Scenario: A few paragraphs describing a situation that professional practitioners in the field might face. A Scenario will deal with an issue relating to the objectives covered in the chapter, and it includes the kinds of details that make a difference.

CASE STUDY: TRUSTWORTHY COMPUTING

ANALYSIS

To its credit, Microsoft has not indicated that this is an easy task that can be solved by a couple of months of code review and programmer training. Further explanation of the long-term (10–15 year) commitment necessary for the success of the vision, and the necessity that all organizations participate, is illustrated in a later whitepaper delivered by Craig Mundie, Senior Vice President and CTO, Advanced Strategies and Policy. You can read this paper at <http://www.microsoft.com/presspass/exec/craig/05-01trustworthywp.asp>.

Many were quick to criticize the memo as just a marketing ploy. Microsoft has been heavily criticized for a long time for producing security-weak, buggy products. This memo was seen as an attempt to change public attitude without doing anything. Microsoft also announced an immediate month-long shut down of work on .NET, the next version of the Windows operating system. The announced purpose was the training of programmers on writing secure code and the scouring of .NET and other existing product code for software bugs. Various sources at Microsoft claim some 9,000 programmers have been trained and that the shutdown lasted for two months.

Additionally numerous bugs have been corrected and the orientation of .NET changed to focus on security versus features.

In response to the original memo and later announcements, a Web site, www.trustworthycomputing.com, put up a page to refer to a www.google.com search page for “Microsoft security or privacy flaw or flaws or hole or holes.” News of this Web page initially dominated the press response to Microsoft's campaign.

In contrast, vendors who have promoted “trusted systems” engineered to deliver security solutions are seizing the opportunity to advertise their solutions. On-board smart card readers in keyboards, and other hardware devices, as well as specialized BIOS-level routines are touted as the answer in the April 4, 2002, article “Signs of Trustworthy Computing,” available at <http://www.wired.com/news/business/0,1367,51521,00.html>.

Trustworthy Computing is a goal that might not be accomplished for many years, if ever. However, there cannot help but be improvements in computer security along the way.

CHAPTER SUMMARY**KEY TERMS**

- Basic input output system (BIOS)
- Blended malware
- Boot sector virus
- Brute-force attack
- Cache
- Centralized controlled computing
- Centralized systems
- Data consistency

Applications can contribute to the security of our computer systems or continue to add additional vulnerabilities to them. The choice is ours. We must scrutinize the applications that will be used on our systems and within our networks, and we must not forget the application development process and its contribution to security or vulnerability. In addition, we should realize the impact of the Internet, or chats, channels, and email as portals for the distribution of malicious applications as well as harmless ones. It is no longer enough to manage the applications that are part of our organizations' business processes. We must realize how easy it is for peripheral code to enter our systems for good or evil.

Key Terms: A list of key terms appears at the end of each chapter. These are terms that you should be sure you know and are comfortable defining and understanding when you go in to take the exam.

Chapter Summary: Before the Apply Your Knowledge section, you will find a chapter summary that wraps up the chapter and reviews what you should have learned.

EXTENSIVE REVIEW AND SELF-TEST OPTIONS

At the end of each chapter, along with some summary elements, you will find a section called “Apply Your Knowledge” that gives you several different methods with which to test your understanding of the material and review what you have learned.

APPLY YOUR KNOWLEDGE**Exercises****8.1 Researching Business Continuity Plans**

The purpose of this exercise is to rate company plans for business continuity.

Estimated Time: 1 hour

1. Take the time to search online for companies or sites that provide information on business continuity or disaster recovery.

Site	Rating	Comments
http://www.springboardhosting.com/products/managed_services/business.php?link=products	Just an ad; not much information	
http://www.disasterrecovery.com/	Contains a lot of information	A very good section on legislation and what is required as far as disaster recovery
http://www.riskconsult.com/home.html	Insurance/risk	Several articles on insurance, risk assessment
http://www.apexdm.com/	Contains just advertising	
www.tbicentral.com	Interesting articles	Must register

Review Questions

1. Where can you obtain information on the potential for specific natural disasters in your location?
2. Why should businesses have a business continuity plan?
3. Explain the difference between DRP and BCP.
4. Why should a business impact assessment be completed?
5. Identify the type of information you would collect from departments to determine whether a particular business process is a critical operation.
6. How do you determine the amount and nature of resources that will be prepared to successfully recover a business process?
7. Why is plan scope important?
8. If e-commerce operations are co-located, is a backup necessary?
9. Should the business recovery plan indicate anything that can be done before the interruption event occurs?
10. What's the difference between a disaster and a business interruption event?

Exercises: These activities provide an opportunity for you to master specific hands-on tasks. Our goal is to increase your proficiency with the product or technology. You must be able to conduct these tasks in order to pass the exam.

Review Questions: These open-ended, short-answer questions allow you to quickly assess your comprehension of what you just read in the chapter. Instead of asking you to choose from a list of options, these questions require you to state the correct answers in your own words. Although you will not experience these kinds of questions on the exam, these questions will indeed test your level of comprehension of key concepts.

xxxiv HOW TO USE THIS BOOK

Exam Questions: These questions reflect the kinds of questions that appear on the actual vendor exam. Use them to become familiar with the exam question formats and to help you determine what you know and what you need to review or study more.

APPLY YOUR KNOWLEDGE

Exam Questions

- A business impact assessment examines business processes to determine which of the following?
 - Which business processes are the most complex
 - Which business processes use computers
 - Which business processes are critical to the organization's survival
 - Whether a business process needs to be a part of the business continuity plan
- A successful test of a business recovery plan has which following result?
 - A pass or fail
 - Demonstrated recovery of data from a backup
 - A visit to the hot site that reveals appropriate equipment is in place and operational
 - Information that can be used to make the plan more effective and knowledge of the readiness of the staff and availability of the equipment necessary
- If a total disaster (the business facility is completely destroyed) occurs, which type of alternative site is best?
 - Hot site
 - Redundant site
 - Warm site
 - Cold site
- Which requirement is most important during the analysis of the impact of business interruption on a particular business process?
 - How large the data file is
 - Current data duplication efforts already in place
 - The amount of money lost for every day of non-operation
 - Whether the operation directly impacts customers
- The first step of any response to a business interruption event should be what?
 - If human life is at risk, evacuate the premises.
 - Call the proper authorities.
 - Secure critical or sensitive data.
 - Determine the source of the problem.
- Business continuity planning is iterative. In which order should events occur?
 - Plan, train, test, revise
 - Plan, test, train, revise
 - Test, train, revise, plan
 - Plan, revise, test, train
- Data management for e-commerce operations might include several functions designed to ensure 24/7 availability. If all of the following are being used, which of them can be eliminated without jeopardizing full data recovery in the event of a disaster?






Answers and Explanations: For each of the Review and Exam questions, you will find thorough explanations located at the end of the section.





Suggested Readings and Resources: The very last element in every chapter is a list of additional resources you can use if you want to go above and beyond certification-level material or if you need to spend more time on a particular subject that you are having trouble understanding.

APPLY YOUR KNOWLEDGE

Suggested Readings and Resources

- Atreya, Hammond, Paine, Starrett, and Wu. *Digital Signatures*. RSA Press, McGraw Hill, 2002.
- Frankel, Sheila. *Demystifying the IPSec Puzzle*. Artech House, 2001.
- Ganapathi, S.J. "Fingerprint Authentication: Shifting the Electronic Security Paradigm." www.scmagazine.com, February, 2002.
- Gove, Ronald A. "Fundamentals of Cryptography and Encryption." In *Handbook of Information Security Management*, edited by Micki Krause and Harold Tipton, Auerbach, 1999.
- Heiser, Jay. "Introduction to Encryption." In *Handbook of Information Security Management, Fourth Edition*, Volume 2, edited by Micki Krause and Harold Tipton, Auerbach, 2001.
- Murray, William Hugh. "Principles and Applications of Key Management." In *Handbook of Information Security Management*, edited by Micki Krause and Harold Tipton, Auerbach, 1999.
- Schneier, Bruce. *Applied Cryptography, Protocols, Algorithms and Source Code in C, Second Edition*. John Wiley and Sons, 1995.
- Schneier, Bruce. *Secrets and Lies, Digital Security in a Networked World*. Wiley, 2000.
- Vallabhaneni, S. Rao. Chapter 5, "Cryptography." In *CISSP Examination Textbooks*, Volume 1. SRV Publications, 2000.
- <http://www.cryptography.com/> (home of Cryptography Research, Inc. It has links to conference papers, articles on protocols, and crypto author sites).





Introduction

The CISSP exam is the premier information security certification. A CISSP is acknowledged by both employers and consultants as a recognition of maturity, experience, and dedication in the information security industry. CISSPs are recognized as having a breadth of security knowledge unparalleled by other certification holders. Ten diverse domains of knowledge are covered on this exam. In addition to passing an exam, the certification requires candidates to have four years of security experience. You should consult the (ISC)² Web site at www.isc2.org for a complete explanation of what acceptable security experience is.

This book is your one-stop shop. Although everything you need to know to pass the exam is in here, you still must meet the experience and ethical requirements set by the exam board. You do not have to take a class in addition to buying this book to pass the exam. However, depending on your personal study habits or learning style, you might benefit from buying this book *and* taking a class. You can locate a class by visiting the (ISC)² Web site (<http://www.isc2.org/cgi/content.cgi?category=15>).

Training guides are meticulously crafted to give you the best possible learning experience for the particular characteristics of the technologies and management skills covered and the actual certification exam. The training guides provide you with the factual knowledge base you need for the exam but then take it to the next level, with case studies, exercises, and exam questions that require you to engage in the analytic thinking that is needed to pass the CISSP exam.

(ISC)², the governing body of the CISSP exam, requires four years of experience in one or more of the 10 domains covered on the exam. A specific definition of exactly what type of experience qualifies can be found on the Web site (“Guidelines for Professional Experience Requirements,” <http://www.isc2.org/cgi-bin/content.cgi?page=167>).

How This Book Helps You

This book takes you on a self-guided tour of all the areas covered by the CISSP exam and teaches you the specific knowledge you need to achieve your certification. The book also contains helpful hints, tips, real-world examples, and exercises, as well as references to additional study materials. Specifically, this book is set up to help you in the following ways:

- ◆ **Organization**—This book is organized by the (ISC)² Common Body of Knowledge (CBK) domains. No official list of exam objectives exists, but the domain definitions provided by the (ISC)² organization have been organized by the authors into helpful objectives. We have also attempted to make the information accessible in the following ways:
 - The full list of domain and compiled objectives is included in this introduction.
 - Each chapter begins with a list of the objectives to be covered.

2 CISSP CERTIFICATION TRAINING GUIDE

- Each chapter also begins with an outline that provides you with an overview of the material and the page numbers where particular topics can be found.
 - The objectives are repeated where the material most directly relevant to it is covered.
- ◆ **Instructional features**—This book has been designed to provide you with multiple ways to learn and reinforce the exam material. Following are some of the helpful methods:
- **Objective explanations**—As mentioned previously, each chapter begins with a list of the objectives covered in the chapter. In addition, immediately following each objective is an explanation of the objective, in a context that defines it meaningfully.
 - **Study strategies**—The beginning of each chapter also includes strategies for approaching the studying and retention of the material in the chapter, particularly as it is addressed on the exam but also in ways that will benefit you on the job.
 - **Review breaks and summaries**—Crucial information is summarized at various points in the book in lists or tables. Each chapter ends with a summary, as well.
 - **Key terms**—A list of key terms appears at the end of each chapter.
 - **Notes**—Notes contain various types of useful or practical information such as tips on technology or administrative practices, historical background on terms and technologies, or side commentary on industry issues.
- **In the Field sidebars**—These relatively extensive discussions cover material that might not be directly relevant to the exam but that is useful as reference material or in everyday practice. In the Field sidebars also provide useful background or contextual information that is necessary for understanding the larger topic under consideration.
 - **Case studies**—Each chapter concludes with a case study. The cases are meant to help you understand the practical applications of the information covered in the chapter.
 - **Step By Steps**—These are hands-on, tutorial instructions that walk you through a particular function relevant to the exam objectives.
 - **Exercises**—Found at the end of the chapters in the “Apply Your Knowledge” section, exercises are performance-based opportunities for you to learn and assess your knowledge.
- ◆ **Extensive practice test options**—The book provides numerous opportunities for you to assess your knowledge and practice for the exam. The practice options include the following:
- **Review questions**—These open-ended questions appear in the “Apply Your Knowledge” section at the end of each chapter. They allow you to quickly assess your comprehension of what you just read in the chapter. Answers to the questions are provided later in a separate section titled “Answers to Review Questions.”
 - **Exam questions**—These questions appear in the “Apply Your Knowledge” section. You can use them to help determine what you know and what you need to review or study further. Answers and explanations for these questions are provided in a separate section titled “Answers to Exam Questions.”

- ◆ **Final Review**—This part of the book provides three valuable tools for preparing for the exam:
 - **Fast Facts**—This condensed version of the information contained in the book is extremely useful for last-minute review.
 - **Study and Exam Day Tips**—You should read this section early on, to help develop study strategies. This section also provides valuable exam-day tips and information on exam/question format.
 - **Practice Exam**—A practice test is included. Questions on this practice exam are written in styles similar to those used on the actual exam. You should use the practice exam to assess your readiness for the real thing. Use the extensive answer explanations to improve your retention and understanding of the material.
 - **PrepLogic**—The *Preview Edition* of the PrepLogic software, which is included on the CD-ROM, provides further practice questions.

NOTE For a description of the *PrepLogic, Preview Edition* software, please see Appendix D, “Using the *PrepLogic, Preview Edition* Software.”

The book includes several other features, such as a section titled “Suggested Readings and Resources” at the end of each chapter that directs you to additional information that can aid you in your exam preparation and real-life work. There are valuable appendixes as well, including a glossary (Appendix A), an overview of the certification process (Appendix B), a description of what is on the CD-ROM (Appendix C), and a discussion of the *PrepLogic, Preview Edition* software (Appendix D).

For more information about the exam or the certification process, refer to the (ISC)² Web site at www.isc2.org.

WHAT THE CISSP EXAM COVERS

The CISSP exam covers a broad range of information security subjects. They are organized into 10 domains. The domains are

- ◆ 1. Access Control Systems and Methodology
- ◆ 2. Telecommunications and Network Security
- ◆ 3. Security Management Practices
- ◆ 4. Application and Systems Development Security
- ◆ 5. Cryptography
- ◆ 6. Security Architecture and Models
- ◆ 7. Operations Security
- ◆ 8. Business Continuity Planning (BCP) and Disaster Recovery Planning (DRP)
- ◆ 9. Law, Investigation, and Ethics
- ◆ 10. Physical Security

Each of these domains is broken down into specific exam objectives. Before taking the exam, you should be proficient in each of the objectives within each domain. These objectives and subobjectives are described in the following sections.

Domain 1: Access Control Systems and Methodology

Discuss the relationship between access control and accountability.

4 CISSP CERTIFICATION TRAINING GUIDE

Define common access control techniques:

- Discretionary access control
- Mandatory access control
- Lattice-based access control
- Rule-based access control
- Role-based access control
- The use of access control lists

Detail the specifics of access control administration.

Explain access control models:

- Biba
- Clark and Wilson
- Non-Inference Model
- State Machine Model
- Access Matrix Model
- Information Flow Model

Explain identification and authentication techniques.

Discuss centralized/decentralized control.

Describe common methods of attack.

Explain intrusion detection.

Domain 2: Network and Telecommunications

Identify the key areas of knowledge of telecommunications and network security.

Explain the International Standards Organization/Open Systems Interconnection (ISO/OSI) layers and characteristics, including

- Physical layer
- Data Link layer

- Network layer
- Transport layer
- Session layer
- Presentation layer
- Application layer

Describe the design and function of communications and network security, including the following:

- Physical media characteristics (such as fiber optics/coaxial/twisted pair)
- Network topologies (for example, star, bus, and ring)
- IPSec authentication and confidentiality
- TCP/IP characteristics and vulnerabilities
- Local area networks (LANs)
- Wide area networks (WANs)
- Remote access/telecommuting techniques
- Secure Remote Procedure Call (S-RPC)
- Remote Access Dial-In User System/Terminal Access Control
- Access system (RADIUS/TACACS)
- Network monitors and packet sniffers

Describe the components, protocols, and services involved in Internet/intranet/extranet design, including the following:

- Firewalls
- Routers
- Switches
- Gateways
- Proxies

- **Protocols**
 - **Transmission Control Protocol/Internet Protocol (TCP/IP)**
 - **Network Layer Security Protocols (IPSec, SKIP, SWIPE)**
 - **Transport Layer Security Protocols (SSL)**
 - **Application Layer Security Protocols (S/MIME, SSL, SET, PEM) (SSL is commonly considered to reside between the Transport and the Session layers.)**
 - **Challenge Handshake Authentication Protocol (CHAP)**
 - **Password Authentication Protocol (PAP)**
 - **Point-to-Point Protocol (PPP)/Serial Line Internet Protocol (SLIP)**
- **Services**
 - **HDLC**
 - **Frame relay**
 - **SDLC**
 - **ISDN**
 - **X.25**

Define and describe communications security techniques to prevent, detect, and correct errors so that integrity, availability, and confidentiality of transactions over networks may be maintained:

- **Tunneling**
- **Virtual private network (VPN)**
- **Network monitors and packet sniffers**
- **Network address translation**
- **Transparency**
- **Ash totals**

- **Record sequence checking**
- **Transmission logging**
- **Transmission error correction**
- **Retransmission controls**

Define and describe specific areas of communication and how they can be secured:

- **Email security**
- **Facsimile security**
- **Secure voice communications**
- **Security boundaries and how to translate security policy to controls**

Explain current forms of network attacks and their countermeasures, including

- **ARP**
- **Brute force**
- **Worms**
- **Flooding**
- **Eavesdropping**
- **Sniffers**
- **Spamming**
- **PBX fraud and abuse**

Domain 3: Security Management and Practices

Understand the principles of security management.

Know what management's responsibility is in the information security environment.

Understand risk management and how to use risk analysis to make information security management decisions.

6 CISSP CERTIFICATION TRAINING GUIDE

Know how to set policies and how to derive standards, guidelines, and implement procedures to meet policy goals.

Set information security roles and responsibilities throughout your organization.

Understand how the various protection mechanisms are used in information security management.

Understand the considerations and criteria for classifying data.

Determine how employment policies and practices are used to enhance information security in your organization.

Use change control to maintain security.

Know what is required for security awareness training.

Domain 4: Applications and Systems Development

Explore software/data issues and describe software and data handling applications. Demonstrate an understanding of the following:

- Challenges of a distributed/nondistributed environment
- Databases and data warehousing issues
- Storage and storage systems
- Knowledge-based systems
- Web services and other examples of edge computing

Discuss the types of attacks made on software vulnerabilities.

Describe and define malicious code.

Discuss system development controls.

Use coding practices that reduce system vulnerability.

Domain 5: Cryptography

Discuss the uses of cryptography including confidentiality, integrity, authentication and nonrepudiation.

Compare and contrast symmetric and asymmetric algorithms.

Describe PKI and key management.

Detail common methods of attacking encryption including general and specific attacks.

Domain 6: Security and Architecture Models

Explain the difference between public versus government requirements for security architecture and models.

Discuss examples of security models including the following:

- Bell-LaPadula
- Biba
- Clark-Wilson
- Access control lists

Explain the basics of security architecture.

Describe and contrast information system security standards including:

- Trusted Computer System Evaluation Criteria (TCSEC)
- Information Technology Security Evaluation Criteria (ITSEC)
- Common Criteria

Describe the Internet Protocol Security (IPSec) standard.

Domain 7: Operations Security

Identify the key roles of operations security:

- Identify resources to be protected.
- Identify privileges to be restricted.
- Identify available controls and their type.
- Describe the OPSEC process.

Define threats and countermeasures.

Explain how audit and monitoring can be used as operations security tools:

- Explain how audit logs can be used to monitor activity and detect intrusions.
- Discuss intrusion detection.
- Explain penetration testing techniques.

Define the role of Administrative management in operations security.

Define operations security concepts and describe operations security best practices:

- Explain antivirus controls and provisions for secure email.
- Explain the purpose of data backup.
- Detail how sensitive information should be handled.
- Describe how media should be handled.

Domain 8: Business Continuity and Disaster Recovery Planning

Explain the difference between disaster recovery planning (DRP) and business continuity planning (BCP) and the importance of developing plans that include both.

Document the natural and man-made events that need to be considered in making disaster recovery and business continuity plans.

Detail the business continuity planning process:

- Explain the process of business impact assessment.
- Define the process of developing the scope of a business continuity plan, including organization analysis, resources, and legal and regulatory requirements.
- Develop business recovery strategies, including planning for crisis management; arranging for cold, hot, warm, and mobile recovery sites; communicating with personnel and management; and developing emergency response and implementation plans.

Detail the disaster recovery planning process, including recovery plan development, implementation, maintenance, and the restoration of business functions:

- Define the process of recovery plan development.
- Describe emergency response, including the development of emergency response teams and procedures. Include disaster recovery crisis management and communication plans.

8 CISSP CERTIFICATION TRAINING GUIDE

- Explain the necessary components of reconstruction procedures, including reconstruction from backup, movement of files from offsite storage, and loading of software, software updates, and data.

Explain the need for, and development of, a backup strategy. Include information on determining what to back up, how often to back up, as well as the proper storage facility for backups.

Domain 9: Law, Investigation, and Ethics

Explain the fundamentals of law.

Define what constitutes a computer crime and how such a crime is proven in court.

Explain the laws of evidence.

Introduce techniques for obtaining and preserving computer evidence.

Identify and plan for computer security incidents.

Discuss computer ethics.

Domain 10: Physical Security

Understand the idea of classifying assets and identifying threats and countermeasures that apply to classes.

Understand some of the most common vulnerabilities and how they affect different asset classes differently. These include

- Understand general principles that apply to the theft of information and assets.

- Know the general criteria that apply to the location and construction of facilities.
- Understand basic methods of controlling physical access to an area.
- Know the basic issues relating to regulating the power supply for computers and other equipment.
- Understand common sources of exposure to water and simple countermeasures.

Understand some of the most common vulnerabilities and how they affect different asset classes differently.

Know the elements involved in choosing, designing, constructing and maintaining a secure site. Elements include

- Site Location and Construction
- Physical Access Controls
- Power
- Environmental Controls
- Water Exposure Problems
- Fire Protection and Prevention

Understand issues and controls related to removable electronic media.

Understand issues relating to storage of paper.

Know the most common issues relating to disposal or erasure of data.

Describe physical intrusion detection methodologies and products.

ADVICE ON TAKING THE EXAM

More extensive tips are found in the “Study and Exam Prep Tips” section, but keep this advice in mind as you study:





- ◆ **Read all the material**—The CISSP domains are broad, and no official list of objectives is published. Instead, any applicant can obtain an (ICS)² study guide that defines the domains and an extensive recommended reading list. You can obtain your copy directly from www.isc2.org. Distributing the guide is not permitted.
- ◆ **Do the Step By Steps and complete the exercises in each chapter**—They will help you clarify the concepts introduced in the text.
- ◆ **Use the exam questions to assess your knowledge**—Don’t just read the chapter content; use the exam questions to find out what you know and what you don’t know. If you are struggling, study some more, review, and then assess your knowledge again.
- ◆ **Review the objectives**—Develop your own questions and examples for each objective listed. If you can develop and answer several questions for each objective, you may find the exam less difficult to pass. If you develop a question for which you can’t find the answer in the book, do go ahead and find the answer elsewhere. The CISSP exam is constantly evolving, and so is the information security profession. This additional knowledge may prove to be valuable, perhaps essential, to you some day.

NOTE

Exam-Taking Advice Although this book is designed to prepare you to take and pass the CISSP certification exam, there are no guarantees. Read this book, and work through the questions and exercises. When you feel confident, take the practice exam and additional exams provided in the PrepLogic test software. Your results should tell you whether you are ready for the real thing.

When taking the actual certification exam, make sure you answer all the questions before your time limit expires. Do not spend too much time on any one question. If you are unsure about the answer to a question, answer it as best you can; then mark it for review when you have finished the rest of the questions.

Remember that the primary object is not to pass the exam, but to understand the material. When you understand the material, passing the exam should be simple. Good luck!







P A R T

I

EXAM PREPARATION

- 1 Access Control Systems and Methodology**
- 2 Telecommunications and Network Security**
- 3 Security Management and Practices**
- 4 Applications and Systems Development Security**
- 5 Cryptography**
- 6 Security Architecture and Models**
- 7 Operations Security**
- 8 Business Continuity Planning and Disaster Recovery Planning**
- 9 Law, Investigation, and Ethics**
- 10 Physical Security**



OBJECTIVES

Discuss the relationship between access control and accountability.

- ▶ With any system, there is information that you want to protect and limit who can gain access to it. Access controls are key to limiting who is allowed to do what on your system. This objective looks at various types of access control and what you can do to protect your system.

Define common access control techniques:

- **Discretionary access control**
 - **Mandatory access control**
 - **Lattice-Based access control**
 - **Rule-Based access control**
 - **Role-Based access control**
 - **The use of access control lists**
- ▶ As with many things, there are many ways to achieve security and many techniques to achieve proper access controls. This objective looks at the various strategies for obtaining an acceptable level of access control across your organization.

Detail the specifics of access control administration.

- ▶ With any organization, there is continual change occurring, and security is continually changing and must be updated periodically. Access control is no exception and must be kept up-to-date and administered on a regular basis.



C H A P T E R

1

Access Control Systems and Methodology

OBJECTIVES

Explain access control models:

- **Biba**
 - **Clark-Wilson**
 - **Non-Inference Model**
 - **State Machine Model**
 - **Access Matrix Model**
 - **Information Flow Model**
- ▶ Throughout the years, many organizations (especially government-based organizations) have developed models to help explain how access control works. This section looks at the various models, including some that have been ported to commercial-based companies.

Explain identification and authentication techniques.

- ▶ To provide proper access controls, the system needs some way to identify who you are and then authenticate you are who you say you are. For example, when you deposit money at a bank they will trust you when you identify yourself. However, when you try to withdraw money, the bank then authenticates that you really are who you say you are by looking at your driver's license.

Discuss centralized/decentralized control.

- ▶ Depending on the size of the organization, there are many ways to manage access control. The two most common approaches are centralized and decentralized controls.

Describe common methods of attack.


- ▶ The best way to understand your risks is to think like an attacker and try to break into your system. This section examines common methods of attacks and what can be done to protect against them.

Explain intrusion detection.

- ▶ The ultimate goal of an attacker is to gain access to a system. The way you gain access is by defeating access controls because access controls are the gatekeepers of your system. By understanding intrusion detection you will gain the ability to protect your access control mechanisms.



OUTLINE

Introduction	17	Access Control Methodologies	37
		Centralized/Remote Authentication	
Accountability	18	Access Controls	38
		Decentralized Access Control	38
Access Control Techniques	19	Domains	39
Discretionary Access Control	20	Trust	40
Mandatory Access Control	21	Methods of Attacks	40
Lattice-Based Access Control	22	Brute-Force	41
Rule-Based Access Control	25	Denial-of-Service	42
Role-Based Access Control	26	Spoofing	42
Access Control Lists	27	Sniffing	43
Access Control Administration	27	Monitoring	43
Account Administration	28	Intrusion Detection	43
		Types of Intrusions	44
Access Control Models	29	Intrusion Prevention	46
Bell-LaPadula	30	How Intrusion Detection Works	46
Simple Security	30	Signature Matching	46
Star Property	31	Anomaly Detection	47
Biba	32	Penetration Testing	48
Summary of BLP and Biba	33	Penetration Testing Versus Security	
Liptner's Lattice	33	Assessments	49
Non-Inference Models	33	Ethical Issues	49
		Performing a Penetration Test	50
Identification and Authentication		Common Tools	51
Techniques	34	Chapter Summary	53
Passwords	35		
One-Time Passwords	35	Apply Your Knowledge	55
Challenge Response	36		
Biometrics	36		
Tickets	36		
Single Sign-On	37		



STUDY STRATEGIES

- ▶ Read each section carefully and make sure you understand the concepts.
- ▶ Apply the concepts that are described in each section to see how they fit or how they could fit into your organization.
- ▶ After you complete the chapter, look at how each of the concepts is interrelated and how together they result in a comprehensive security solution.



“Access control is the collection of mechanisms that permits managers of a system to exercise a directing or restraining influence over the behavior, use, and content of a system.

The candidate should fully understand access control concepts, methodologies, and implementation within centralized and decentralized environments across the enterprise’s computer systems. Access control techniques, detective and corrective measures should be studied to understand the potential risks, vulnerabilities, and exposures.”

—Common Body of Knowledge study guide

This chapter covers Domain 1, Access Control Systems and Methodology, which is 1 of 10 domains of the Common Body of Knowledge (CBK) covered in the Certified Information Systems Security Professional Examination. This domain is divided into several objectives for study.

INTRODUCTION

A key part of security is controlling access to critical information. This chapter examines the various schemes used for accomplishing this. In talking about access control, it is important that we distinguish between *authentication* and *access control*. Some people think of the two terms as being similar or interchangeable, but they are quite different. Passwords and similar techniques usually provide only authentication—they identify a user and verify that the user is who he says he is. Just because you know that a certain person is actually Bob does not mean that Bob should have access to every piece of data on your network. That is where access controls come in. After you properly identify a user, you then want to control what access he is given on the system. In most cases, you want to give the user the least amount of access he needs to do his job and nothing else. This concept is often referred to as the *principle of least privilege*. It gives you the power of combining authentication with access control.

Both authentication and access control are needed to achieve a high level of security. One without the other leaves huge security holes that allow an attacker a high chance of compromising a target’s network.

NOTE

Legislation of Privacy and Security

Because privacy of personal data and the security of the systems that contain this and other sensitive information are of increasing concern, legislation has been written to address the issue. The Health Insurance Portability and Accountability Act of 1996 (HIPPA, <http://cms.hhs.gov/hipaa/>) dictates how patient data should be protected at hospitals, insurance companies, and other places it might be collected and used. The Gramm-Leach-Bliley Act includes regulations that "...require clear disclosure by financial institutions of their privacy policy regarding the sharing of non-public personal information with both affiliates and third parties." See <http://www.senate.gov/banking/conf/grmleach.htm> for more information.

For example, if you have only authentication and no access controls in place, it might take a while for an attacker to compromise an account and guess a password, but once he does, he has full access to the systems. When no access controls are in place, there is nothing stopping anyone from getting to any piece of data that he wants. Also, by having only authentication, an internal user is allowed full access, which could cause a lot of damage either intentionally or accidentally. I have been involved with more incidents that were caused by accidents because legitimate users had more access than they should have had and accidentally caused major network problems.

Looking at it from the other perspective, having access control with no authentication means that people are limited to what they can do on your network. However, because you have no way to identify a given user, anyone could impersonate any other user to get the access he needs. So, even though Bob has limited access, he could impersonate the root account—which has full access—and do whatever he wants on the system. Nothing would stop Bob from doing this because no authentication is done against anyone, so the system believes whatever the user tells it. This, as you can imagine, is extremely dangerous and hardly ever done. It is more common to see authentication without access controls, rather than access controls with no authentication.

ACCOUNTABILITY

Discuss the relationship between access control and accountability.

Would anyone follow the speed limit if we knew for a fact that there was no chance we would be pulled over? If there is no chance that we could be held accountable for our actions, there is a good chance that most people would drive as fast as they possibly could. Now, there are certain people, like my dad, who would drive 55mph no matter what the speed limit was, but most of us are kept honest because we know there is a chance that there could be a cop around any corner who would hold us accountable by giving us a ticket. This same concept of accountability is critical when it comes to security—mainly access controls.

Even if you control who can do what on a system, you want to be able to track this information to hold people accountable for their actions. Also, in some cases the access controls might not have been set up correctly. You might have given a user too much or too little access. By holding a user accountable, you can see exactly what she did or did not do and use this information to adjust the access controls to the proper level. Remember: Security is a never-ending job. Just because your access controls are correct today does not mean they will be correct tomorrow.

A common way to keep track of accountability is with *logging*. By recording what people do on a system, you can hold them accountable for their actions. Actually, there is one other piece that is missing for accountability to work—you have to know that a one-to-one relationship exists between an account/user ID and an individual. If you cannot prove to a reasonable level that Bob is the only one who should have access to Bob's password and therefore be the only one logging in with the Bob account, accountability does not work. The biggest problem with accountability is shared accounts. Except in certain extreme circumstances, shared accounts must be avoided. This policy must be clearly reflected in the security policy and strictly enforced; you must not tell anyone else the password for your individual account under any circumstances. In cases where someone forgets her password and the help desk has to change her password for her, the first time she logs on, she must be forced to change her password. This way, no one else can log in to the system as a different person and impersonate another employee. You must make people accountable for their actions so you can properly enforce access controls.

ACCESS CONTROL TECHNIQUES

Define common access control techniques:

- **Discretionary access control**
- **Mandatory access control**
- **Lattice-Based access control**
- **Rule-Based access control**
- **Role-Based access control**
- **The use of access control lists**

Access controls are important, but how do you determine the proper access controls an individual or entity should have on a system. There are two general types of access control: discretionary and mandatory access controls. They are often referred to by their acronyms: DAC (discretionary access control) and MAC (mandatory access control).

Discretionary Access Control

Discretionary access control is essentially based on human decisions about whether someone (or a service, an application, and so on) should be allowed access to a particular resource, such as a file or directory. Most companies implement this across their organizations. They might have guidelines or policies that say if you work in this department, you can access only these directories and the people who set up these accounts religiously follow these policies. The problem with DACs is that they are controlled by humans, which means they are open to mistakes and can easily be overwritten.

For example, when an administrator is adding a new account, he might accidentally give a user more access than she should have because he erroneously thought the user worked in a different department. This is less of a problem with DAC than humans overriding the access controls.

In an organization, when a certain individual wants access to something he does not have proper permissions for, he usually whines and kicks and screams like a 2-year-old having a temper tantrum. He keeps going up the chain of command, complaining that he cannot get his job done because he does not have access. Eventually, a higher-level manager calls the person responsible for setting the permissions and tells her to give the user permissions because the manager is tired of hearing him whine.

For these reasons, DAC does not provide a high level of access control because the measure of who should have access is very subjective. A human can give and take controls based on her mood, who her friends are, or who yells at her. This is okay for some environments, but for other environments a higher level of protection is needed.

Mandatory Access Control

Mandatory access control applies a higher level of access control in which the computer system strictly controls who can access what resources. Because MAC is based on using classification levels, it is more popular in government-type environments. However, it is slowly working its way into the commercial arena and is starting to show up in financial institutions. Every entity using the system gets a classification level. So, each user has a classification level associated with her account and each piece of data has a classification level. When a user tries to access a piece of data, the system determines whether she can access that piece of data by looking at both the classification of the user and the classification of the data.

With MAC it is important to highlight a couple of key points. First, users could have multiple accounts associated with different levels of access. For example, Bob could have a secret account and a top secret account, and depending on the work he is trying to do, he would log in to the appropriate account. Some people might ask why Bob would do that when he could just always log in as the top secret account and access everything instead of switching between accounts. This logic makes an assumption that one level of access encompasses the level below it. For example, we are all familiar with the government model in which unclassified is the lowest level. The next level is confidential, so someone with this access can access anything labeled unclassified and confidential. At the secret level you can access unclassified, confidential, and secret information. In this case, you have to trust the user to log in to the account with the least access he needs to do his job. You can quickly see this as a limitation of MAC.

Another possible alternative is to use access levels that are not all inclusive. So, one access level does not mean you can access anything in a lower level because the levels are not set up in a hierarchical fashion. This is often called *compartmentation*. Think of a typical corporation. You might have a finance compartment, an HR compartment, and an engineering compartment, just to name a few. The director of engineering might have two accounts: one with HR access when he is hiring or firing people and one for engineering when he is looking at the progress of certain projects. The director would log in with the proper account based on the type of work he is going to perform.

What usually occurs with MAC is that both the hierarchical levels, such as secret and top secret, are combined with compartmentation to provide a finer granularity of control. When the system enforces MAC, it first makes sure you have a level equal to or greater than the data you are trying to access and that you have all the proper compartmentations to access the data. For example, if Bob has top secret access with HR and engineering compartments, he can access data at the secret level with no compartments. He can also access secret data with an HR compartment. However, if Bob tries to access a system at the secret level with a finance compartment, the system will not let him have access. The level of access is appropriate, but because he is missing a compartment, the system denies him access to the data.

One other key point is that when we think of MAC, we are so accustomed to government organizations that we immediately think secret, top secret, and so on, but that does not have to be the case. You can come up with whatever levels of access you want. For example, you could have company proprietary, company sensitive, and executive staff only. These would roughly be equivalent to confidential, secret, and top secret, respectively.

Lattice-Based Access Control

Lattice-based access control is a form of MAC for strictly implementing access controls across an organization. Once again, this model tends to be used in more government-type settings but could also be implemented across a commercial enterprise. With a lattice model, you first have to define a set of security classes that can be assigned to users or objects. For example, a security class could consist of confidential, secret, and top secret. After you have a defined set of security classes, you define a set of flow operations showing when information can flow from one class to another. This is generally depicted as an arrow. So, if you have this:

confidential → secret

This means that information can flow from confidential to secret. By careful examination, you can also see that because there is no flow operator or arrow from secret to confidential, information cannot flow from secret to confidential. Remember that flow relations are only one way.

Chapter 1 ACCESS CONTROL SYSTEMS AND METHODOLOGY 23

If you want information to flow both ways, you have to explicitly put two arrows in place. For example, if you wanted to have a two-way flow relation, you would write the following:

confidential \rightarrow confidential
confidential \leftarrow confidential

These two statements show that confidential can flow to confidential in either direction. That might seem like a very obvious point, which it is, but it was used to emphasize a point.

Now that we have defined security classes and flow operations, the following are the requirements for a lattice:

- ◆ The security class must be finite and not change.
- ◆ All the flow operations must make a partial order. A partial order has the following properties:
 - **Reflexive**—If you take an item in the security class information, it can always flow back to that same security class. Confidential \rightarrow confidential is an example of the reflexive property.
 - **Anti-symmetric**—If information flows in one direction, it cannot flow back in the opposite direction. For example, note the following:

confidential \rightarrow secret

This means you cannot have the following:

secret \rightarrow confidential

Another way to look at this is that information flow cannot be symmetric; it can flow in only one direction.

- **Transitive**—If information can flow from one security class to another security class by going through a third security class, information can also flow directly between those two security classes. Transitive is easy to see with an example. Note the following:

confidential \rightarrow secret

secret \rightarrow top secret

By the transitive rule, you also must have the following:

confidential \rightarrow top secret

This is the case because, if confidential can flow to secret and secret can flow to top secret, it is also implied that confidential can directly go to top secret, so that information flow must be added.

- ◆ The lattice must have a lower bound, which is usually considered the null class. For example, unclassified could be considered the lower bound because it is the base denominator in which you cannot go any lower.
- ◆ The lattice must have an upper bound, which represents a combination of all the items in the security class.




FIGURE 1.1
A simple lattice-based access control model.




FIGURE 1.2
Lattice shown with reflexive and transitive edges added.

Because lattice-based access controls are usually drawn as directed graphs, a lattice is considered a graph that follows the previous set of rules. Let's look at a simple example of a lattice to emphasize what we mean by lattice-based access controls. Figure 1.1 illustrates the concept of compartments within an organization. Let's say that there is a finance compartment and an engineer compartment. For this example, the security class would consist of two elements (finance, engineer).

Let's go through the four properties to make sure Figure 1.1 is a lattice. The first property is that the security class must be finite and not change. In this case, the security class consists of only two elements—finance and engineer. The second property says it must be a partial order, which implies reflexive, anti-symmetric, and transitive. Typically, when drawing a lattice, you do not draw the reflexive or transitive arrows because they clutter the diagram, but they are implicitly implied by the model. So, information can flow from finance to finance even though it is not explicitly shown with an arrow. Therefore, the reflexive property is true for this diagram. Information can also flow from {} to {finance, engineer}, so the transitive property also holds true.

For it to be a partial order, the last property we have to look at is anti-symmetric. In this case, all the information flows are one way, so the anti-symmetric property is true—meaning this lattice is a partial order. For clarification, the diagram could also be drawn with the reflexive and transitive arrows explicitly added, as shown in Figure 1.2.

Next, you have to ensure that there is a lower bound or a null set from which everything else is derived. In this case, because the {} contains nothing, this is the lower bound. The final criteria is that an upper bound exists, composed of all the elements in the security class. Because only two items are in the security class, the upper bound {finance, engineer} contains them both. Therefore, the diagram is a lattice and does indeed enforce information for using lattice-based access controls.

For those of us who have worked in government facilities, lattice-based access controls might seem easy to understand, but it's important to understand the mechanics behind it.

Rule-Based Access Control

Rule-based access control involves setting up parameters around which an individual can access a system. Usually these parameters are written as rules. Simple rules can be listed as “user Bob can access resource X, but he cannot access resource Y.” They can also become more complex; for example, “user Bob can only access resource X if he is coming from workstation alpha and it is between the hours of 8 a.m. and 5 p.m.”

One of the main reasons rule-based access controls are not very popular has to do with scalability and maintainability. If you have a small organization, keeping a set of rules for each user is manageable. On the other hand, if you have a small organization, you probably do not need such rules because everyone knows his role and is trusted to some extent. Also, if you have a small number of users, you likely have a minimal number of systems, so such controls are not needed.

For larger organizations, such rules could be helpful to truly enforce a principle of least privilege across an organization. This principle says that you should give an entity the least amount of access it needs to do its job and nothing else. With rule-based access controls, you can control with a fine level of granularity who can do what on the system. The first problem with this is gathering the information. Figuring out who should do what and then entering it into the system can be extremely time-consuming. The second problem is having to maintain such a complex list. As people change jobs and transfer responsibility, having to constantly update and correctly maintain this information can be overwhelming even for a large staff of people. That is one reason some companies prefer *role-based access control*.

Role-Based Access Control

Unlike rule-based access control where you give access to individual users, in *role-based access control* you develop roles or positions across your company and assign access to the role based on the job functions of that position. This is the most widely used form of access controls. For example, you might create roles of a junior Windows NT administrator, a mid-level network operator, and a senior-level data center engineer. Because each of these positions has a set job function, they can be given the proper level of access—or the minimal amount of access those positions need to do their jobs and nothing else. After the positions have been defined, you assign people to those positions. When a person is assigned to a given position, he inherits all the permissions or access rights associated with that position.

This approach is much easier to maintain and manage. First, because fewer positions exist in an organization, less work is involved to set up the access permissions. Because people change jobs frequently, when someone moves to a new position, he is removed from the one role and put in a different role. Consequently, he instantly inherits all the proper access needed to do the new job. If a new position is created, a new profile has to be created with the proper positions.

The real power of role-based access control is when you have to change the permissions associated with a given role. Let's say that a position of senior network engineer has 30 people with that job function. Without role-based access control, if the function of that job changes, 30 different people would have to have their access individually tracked and changed. As you can imagine, this would involve a lot of work. With role-based access control, you just have to change the access associated with the role and perform that once, and all 30 people would automatically be updated with the new access they need.

Role-based access control is typically implemented by using *groups*. You create a group and give permissions to that group. User accounts are then added to groups based on job function. When a user switches positions, he is removed from one group and added to another group. In practice, using groups to implement role-based access controls is usually a little complicated. The reason is that not everyone in a given position requires the exact same level of access. For example, a senior network engineer has a wide range of responsibilities, and not every senior engineer performs the same functions.

By creating a single group that has all the possible access a senior engineer might need, some people have more access than what they need to perform their jobs. This breaches the principle of least privilege. Therefore, groups are typically created based on certain levels of functionality; then, a given position might have three to four groups associated with it. When a person is given a new role or position based on what functions he will perform, he is added to the appropriate groups.

Access Control Lists

Access controls lists (ACLs) are similar to rule-based access control but tend to be more formalized. With ACLs, you create a list of rules usually based on IP addresses or some piece of information that can easily be discernible in the packets that go across the network. For each rule, you specify whether you will allow or deny traffic. ACLs are often associated with routers and applied to limit the amount of traffic that can go to a given network resource.

ACLs are often implemented at border routers to provide a very basic level of access control. The most popular ACLs are used on Cisco routers. The following is a basic ACL:

```
Access-list 1 deny 10.0.0.0 0.255.255.255
```

```
Access-list 1 permit any
```

```
Access-list 1 deny 0.0.0.0 255.255.255.255
```

This denies access from the 10.x.x.x network and allows any other traffic. Essentially, any IP address whose first octet is 10 is denied access, but any other IP address is allowed or permitted.

ACCESS CONTROL ADMINISTRATION

Detail the specifics of access control administration.

As with any aspect of security, setting it up is not the difficult part—it is the ongoing maintenance and enforcement that is the most difficult. Access control is no exception.

Access control essentially involves two pieces of information—a user ID and a password. This information must be set up and maintained for each user of the system. When a new employee starts at the company, she must be added in a timely fashion, and when someone leaves the company, the account must be disabled in just as timely a fashion.

IN THE FIELD

DISABLING VERSUS DELETING

Notice the key word when someone leaves the company—you *disable* her account; you do not *delete* her account. It is a common mistake to delete accounts when people leave the company. Instead, you should disable the account for a certain period of time. Then, after an account has been disabled for a certain period, you can delete it. This is done for two main reasons. First, it is common for people to leave a company or think they are leaving a company and then decide to come back to work for that same company. Second, some operating systems remove access to resources when you delete an account. If a company has a marketing employee who has left the company and she is being replaced by a new employee, you want the new employee to have the same access as the old employee. If the old employee's account was deleted, you have no idea what access she had. So, assigning access to the new employee is more difficult. On the other hand, if you just disabled the old employee's account, you could rename it to the new employee so he instantly has all the same access the previous employee had.

Account Administration

When a new account is set up, the administrator needs to assign a temporary password for the account. It is recommended that you create an initial random password for each account as opposed to using a standard account across a company. If a standard password is used across a company then whenever a new account is created or a password is reset on the account, anyone who knows the standard password could get access to the account. It is better to generate a unique password for each account; then when the user needs to log on, she can call the help desk to get the new password.

The first time the person logs in with the temporary password she is forced to change her password to something that only she knows. Access control works only if a single person is the only one who has access to a given account or is the only one who knows the password. If multiple people have access to the same password, you lose accountability for who is doing what on your systems and network. Keeping a one-to-one relationship between accounts and employees is an easy way to track who is doing what. You monitor and keep track of access controls through *logging*. It is recommended when logging events to log both success and failures. Some administrators log only failures, but this does not give you sufficient information to make decisions. For example, if you logged only failed events, you would not have the complete picture of what is happening on your network. Let's say your logs show five failed logon attempts for Sally, followed by five failed logon attempts for Bob. You know that someone is trying to gain access, but you do not know whether he actually got into Sally's account or whether he got tired and moved on to Bob's. Only by showing both failed and successful attempts can you tell whether someone actually gained access to a given account.

When assigning permissions to accounts, you should give someone the least amount of access he needs to do his job, and nothing else. Notice that you should give people enough access to do their jobs and take away all other extraneous access to this system.

Also, for access to sensitive information, you should maintain a *separation of duties*. This involves taking sensitive access and breaking it up among several individuals. If access is needed to this information, multiple people must participate to gain access. This is often seen in military movies where access is needed to nuclear weapons. Two people must both insert their keys and turn them at the same time to get the necessary access.

ACCESS CONTROL MODELS

Explain access control models.

This section covers some strategies or models for implementing access controls across an organization. These models serve as rules for the road when figuring out some general principles that should be followed when implementing access control.

Most of these were originally developed with a government slant but can easily be applied to commercial settings. However, the examples in this section use the general government classification scheme of unclassified, confidential, secret, and top secret, where unclassified is the lowest and top secret is the highest. The reason this is done is that even people who have not worked for the government understand this hierarchical scheme, which makes explaining the topics easier.

The following are the models discussed:

- ◆ Bell-LaPadula
- ◆ Biba
- ◆ Liptner
- ◆ Non-inference

Bell-LaPadula

The Bell-LaPadula (BLP) model deals with the flow of information from a confidential standpoint. Remember that the definition of confidential is to prevent, detect, and deter unauthorized access to information. This is used when you have a secret and do not want someone else to be able to read it. The BLP protects people from accessing information they should not have access to. BLP is composed of two rules:

- ◆ Simple security deals with reading information or files.
- ◆ The star property deals with writing information or creating new files.

Simple Security

The simple security rule deals with reading information and ensures that someone cannot read information they do not have access to read.

The simple security rule states that a principal (P) can read an object (O) only if the security label of P is higher than (or equal to) the security level of O. This means that information can flow from security level O to security level P. An example might help explain this:

NOTE

Subject and Principals Before we cover the two rules, there is a concept of subject/principal versus users when we talk about BLP. In BLP these rules apply to subjects or principals—people who have normal access to the system. In the BLP model users are considered trusted entities and will not disclose information outside the computer system; principals are not considered trusted.

If a principal has secret clearance and he wants to read an object that is labeled as top secret, the request will not be allowed because the object has a higher clearance than what the user has. It makes sense that someone with secret clearance cannot access top-secret information. However, the principal who has secret access can read an object with a secret, confidential, or unclassified security label. This is because those security labels are equal to or lower than the security label the user possesses.

Star Property

The star, or *, property deals with the writing of information. It states that a principal (P) can write to an object (O) only if the security label of O is higher than (or equal to) the security label of P. This means information can flow from security label P to security label O. This rule states that a user can write to an object only if the security label is equal to or greater than his own. If a principal has a security label of secret, he can write to an object with a security label of secret or top secret but cannot write to an object with a label of confidential or unclassified. This might seem a little strange, but it is meant to prevent the leakage of information.

The star property is meant to protect against write-down Trojan horses. Let's say that a principal with a confidential security label wants to read a secret document, but the system does not allow him. Someone could insert a Trojan horse into a program that a principal who has a secret security label uses. When he does his work, this Trojan horse works in the background, reads the secret document, and writes it to a confidential document. The evildoer who had only confidential access could now read the information because the Trojan horse put the information in a document with a security label that the principal could access. The star property prevents this from happening.

However, this property is still a little dangerous because it allows a principal to write to a higher level, which could result in an integrity problem. Let's say that a principal has a secret security label and a document is labeled top secret. Even though the principal cannot read the document, he can still write to the document—despite the fact that he does not know what it says. So, this principal could overwrite critical pieces of the document, making the document no longer accurate and resulting in an integrity problem. The principal could also overwrite all the information so no one can read it, resulting in a denial-of-service attack.

In practice, principals are usually allowed to write only to an object that has the same security label. This prevents the write-down Trojan horse and the write-up integrity problems discussed here.

Biba

The Biba model is similar to BLP except for the fact that, instead of dealing with confidentiality, it deals with integrity. It does not care whether someone can gain access to information she should not have access to as long as she cannot change the content so that it is no longer accurate. Biba has the same two rules BLP has:

- ◆ Simple security deals with reading.
- ◆ The star property deals with writing.

The big difference, which seems confusing at first, is that both rules are the opposite of the BLP model.

With BLP, the rule is that you cannot read up—a principal cannot read an object that has a higher security label. Because Biba deals with integrity, the rule is switched to not read down. The simple security rule with Biba says that a principal (P) can read an object (O) only if the security label of O is higher than the security label of P.

The star property of Biba says you cannot write up, which once again is the opposite of BLP. The star property with Biba says a principal (P) can write to an object (O) if the security label of P is higher than the security label of O.

If you examine both models, they are equivalent except for the fact that BLP is a bottom-up model, which says that information can flow from the bottom to the top. Biba, on the other hand, is a top-down model, which means information can flow from the top down.

Summary of BLP and Biba

I would recommend remembering the following key points about BLP and Biba:

BLP model:

- ▶ Simple security
- ▶ Simple property
- ▶ Deals with confidentiality

Biba model:

- ▶ Deals with integrity.
 - ▶ The rules of Biba are the opposite of BLP.
-
-

Liptner's Lattice

As stated earlier, most of the models we have talked about relate to government settings. These models, however, can easily be applied to commercial settings, and that is exactly what Liptner did. He applied lattices and the principles we talked about to non-military examples. Essentially, he changed the labels from terms such as *confidential* and *secret* to *system programmers*, *production code*, and so on.

Non-Inference Models

Non-inference models deal with examining the input to and output from a system and seeing whether you can infer any information that you should not have access to. These models tend to be more theoretical in nature, but they are still beneficial to understand. The general principle is that you have a system with several inputs and several outputs, and if you modify or purge any of the inputs, the outputs should remain unchanged. The reason for this is if you can modify an input and a one-to-one relationship exists between inputs and outputs, an output would change and you could start to infer information about the system.

IDENTIFICATION AND AUTHENTICATION TECHNIQUES

Explain identification and authentication techniques.

From an access control standpoint, you have to tell the system who you are and then prove to the system you are who you say you are. For example, when I go to the airport to pick up my electronic ticket, I walk up to the counter and say my name and where I am going. Based solely on that, they look up my information and find my reservation. However, before they will give me the ticket, I have to prove to them I am really who I say I am—and I usually do that with either a driver's license or a passport. After I have authenticated, they give me my ticket. The same thing has to be done when you try to gain access to a computer system. You present a user ID and then a password to gain access. The user ID usually consists of some combination of the user's first name and last name. Even though the user ID is not meant to be secure, if someone is able to guess a user ID, gaining access is slightly easier. After an attacker knows she has a valid user ID, it is just a matter of guessing passwords to try and get in. Hopefully, that will be impossible because everyone has a very strong password. Unfortunately, that is not reality. Because people tend to pick weak passwords, figuring out valid user IDs gives attackers the edge. Ideally, using user IDs that are not predictable makes an attacker's job that much more difficult.

In terms of proving you are who you say you are, there are several techniques for doing this:

- ◆ Passwords
- ◆ One-time passwords
- ◆ Challenge response
- ◆ Biometrics
- ◆ Tickets
- ◆ Single sign-on (SSO)

Three things can be used to authenticate yourself:

- ◆ Something you know—passwords
- ◆ Something you have—one-time passwords
- ◆ Something you are—biometrics

These are discussed in the following sections.

Passwords

A *password* is typically a word the user picks to prove he is the owner of the account. The problem with typical passwords is users tend to choose easy-to-guess passwords. Even with password policies, users still pick passwords that are composed of dictionary words because they're easy to remember. It is recommended that you encourage users to pick passwords that are long; contain lowercase letters, uppercase letters, numbers, and special characters intermixed; and contain no dictionary words within.

Because users are ultimately in control of what passwords they choose, authentication methods based on user-derived passwords tend to be weak. Even if a company automatically derives the password for the user, this is still not considered strong because the password is now hard to remember, so most people will write it down. This defeats the purpose of having a strong password.

One-Time Passwords

One-time passwords solve the problems of user-derived passwords. With one-time passwords, each time the user tries to log on she is given a new password. Even if an attacker intercepts the password, he will not be able to use it to gain access because it is good for only one session. One-time passwords typically use a small hardware device (key fob or SecureID) that generates a new password every minute. The server also has the same software running, so when a user types in her password (off the device), the server can confirm whether it is the correct password. Each time the user logs on she has a new password, so it is much more secure. The problem, however, is that users have to ensure they have the device with them at all times; otherwise, they cannot log on.

In addition, software-based one-time password programs are available, such as S/Key, but they are not as popular as their hardware counterparts.

Challenge Response

An alternative to one-time passwords is *challenge response schemes*. Instead of having the device just blindly generate a password, a user identifies himself to the server, usually by presenting his user ID. The server then responds with a challenge, which is usually a short phrase of letters and numbers. The user types the challenge into the device and, based on the challenge, the device responds with an output. The user then types that output in as his password to the server. This scheme is slightly more complicated, but it allows the password to be based on changing input rather than just time. Also, because the input is not based on time, you do not have to worry about clock skew problems, which happen with one-time passwords. If the clock on the server or the device slowly gets out of sync, eventually the user will be unable to log on to the system.

Biometrics

Both one-time passwords and challenge response schemes have the problem that the user has to carry a device around with him and if he loses the device, he can no longer log on to the system. Biometrics authentication is based on something you are, so you do not have to worry about forgetting a password or leaving a device at home. Several types of biometric devices are available, some of which can be used to authenticate fingerprints and hand, face, and retinal scans. Biometrics are covered in detail under the Physical domain and are mentioned here just for completeness.

Tickets

Another way to authenticate is for the system to give you a *ticket*, and if you can unencrypt the ticket, you can gain access. These schemes rely on the exchanging of keys prior to authentication.

One of the common programs that does this is Kerberos. Before you can use Kerberos, you must exchange a secret key with the server. Only you and the server know the key. When you connect to the system, you just tell the server your user ID, and the server sends back an encrypted ticket. If you're who you say you are, you will know the key and be able to unencrypt the ticket, thereby gaining access to the information; otherwise, you will be denied access. Ticket schemes do not scale very well, which is why they are less common than the other approaches.

Single Sign-On

Single sign-on (SSO) is another scheme for authentication when you have a large number of applications that all need to authenticate the same user. Instead of requiring the user to log on multiple times, she logs on once to a central server and that server authenticates her to the other applications automatically. This lessens the burden on the user because she logs on only once, but it increases the overall security. With SSO, if someone is able to compromise someone else's information, he can gain access to everything. Also, if someone stays logged on and forgets to lock her workstation when she walks away, anyone sitting down at her workstation can have full access to everything without ever having to provide a password. SSO shows the balance that you need to achieve between security and functionality.

ACCESS CONTROL METHODOLOGIES

Discuss centralized /decentralized control.

This section examines two primary remote access controls: Remote Authentication Dial-In User Service (RADIUS) and Terminal Access Controller Access Control System (TACACS). (Actually, in most cases when you see a reference to TACACS, they are referring to TACACS+, which has some more advanced features.)

Centralized/Remote Authentication Access Controls

RADIUS and TACACS+ are usually used interchangeably for remote access controls. They are typically used when users are required to authenticate to different applications and you do not want to manage a separate listing of user accounts for each application. Instead, you would point all the applications to your RADIUS or TACACS+ server to authenticate the users. This way, you have to administer and manage only one set of accounts and credentials.

Another area where you would use RADIUS or TACACS+ is when you have an application or a device that needs to authenticate users but no built-in facility exists for doing this. A good example of this is Cisco routers. The key thing to remember is that if you want to have a centralized access control server for authentication and authorization, RADIUS and TACACS+ provide the facility for doing this. Most of the time, Cisco recommends using TACACS+ with its routers and devices.

Decentralized Access Control

If you want to get into a heated argument with a security professional, just bring up centralized versus decentralized access control. It seems there is no right answer; there is just an answer that depends on the environment you seek to control. With *centralized control*, a single authority or system is responsible for access control. The biggest problem with this is that a single point of failure exists that could also become a bottleneck for an organization. For a small organization, centralized control might make sense, but for a larger organization, this is not practical. One way around this is to implement a centralized model with backup or failover capability. Therefore, even though a single source is managing it, there are several systems, so if one fails, the other one can kick in and take over. Another way around this problem is to use *decentralized control*.

With decentralized control, each individual or department is responsible for its own access control. In the early days of networking this was typically used. For example, with Windows for Workgroups, you set up a network of computers with decentralized access control.

Each user who connected to the network was responsible for setting up access controls for her resources. Essentially, if someone acquired access to your resources and he was not supposed to, it was your fault for not setting up the access controls properly. This, however, has its own set of problems because now you are trusting that each entity responsible for access control does the right thing. For a large organization, this can be a very scary proposition.

In reality, what happens in most situations is a compromise—and access control is no exception. Most organizations tend to use a hybrid. Depending on the size and structure of the organization, they might set up several zones or domains (each with a centralized access control for that domain). Then, to allow each of the domains to access resources in the other domains, they set up trust relationships between the two domains. Let's look at domains and trust in a little more detail.

Domains

A *domain* in its most basic form is a group of computers under the same administrative authority. It is a way to group systems together to make them easier to maintain and control. From an access control standpoint, a domain is a group of systems that all authenticate to a central system or group of systems.

A domain is modeled after the centralized access control model. You usually have several domain controllers that can authenticate users to the network and authorize them to access resources. This way, if one system goes down, it does not present a single point of failure. If each domain controller maintained its own database, things would quickly get out of sync and very messy, so instead a single primary domain controller maintains the master copy of users and passwords. Other domain controllers can authenticate users, but any changes to accounts must be done against the primary domain controller. You might think that by doing this, you create a single point of failure. You do, and you don't. The primary domain controller at regular intervals pushes the new copy of the database to the other domain controllers so they have an updated copy of the information. If the primary domain controller crashes, one of the other domain controllers can take over its role because they have a fairly accurate copy of all the users. This is not perfect because, if a new user is added to the primary domain controller and it immediately goes down before the latest copy is pushed to the other domain controllers, that user is lost.

NOTE

Comparison to Microsoft One general note of caution: Even though Microsoft uses the terms *primary* and *backup* domain controllers, this discussion was written independent of any operating system and is meant to show how the concept of domain fits into access control.

Domains might seem okay, but what happens when you have a user in one domain who wants to access a server in another domain? This is where trust comes into play.

Trust

If your organization sets up a hybrid model with a bunch of domains, the question arises, “How does a user in one domain access resources in another domain?” You do this by setting up *trust relationships* between domains. For example, if you set up a full trust relationship between domain A and domain B, anything in one domain can access something in the other domain.

What does setting up a trust really mean? A trust says that you trust one domain to provide the same level of access control that another domain does. So, if it authenticates a user and thinks that user is worthy enough to access data on her own domain, she can access data on your domain. It is similar to when you go on vacation and leave a key with your neighbor. You would only give the key to a neighbor you trust. This means you expect the neighbor to protect your house just like he would his own house.

When we talk about trust relationships, we sometimes talk about a *full trust* or a *one-way trust*. For example, with the full trust relationship between domain A and B, A’s users can access B’s data and B’s users can access A’s data. However, sometimes you might want to only set up a one-way trust. I have some neighbors who trust me, but I do not trust them. I have a key to their house, but they do not have a key to my house. This is considered a one-way trust, and a similar thing can be done with domains. With a one-way trust, A might trust B but not the other way around, so you set up a one-way trust from domain B to domain A. This says domain A trusts domain B, but domain B does not trust domain A. So, domain B’s users can access domain A’s data, but domain A’s users cannot access domain B’s data. As you can see, you can get very creative with trust relationships.

METHODS OF ATTACKS

Describe common methods of attack.

The only way to have a good defense is to understand the offense and know how it operates. Access control is no exception. To ensure that you have proper access control and that it is set up correctly, you need to understand how attackers try to break access controls. By understanding how someone is trying to break them, you can build better defenses that either eliminate the threat or make it much harder for the attacker to succeed.

Types of attacks include

- ◆ Brute-force
- ◆ Denial-of-service
- ◆ Spoofing
- ◆ Sniffing

Each of these is discussed in the following sections.

Brute-Force

With a *brute-force attack*, an intruder tries all possible combinations until she guesses the right one. Brute-force attacks are most popular with cracking passwords. A lot of people do not realize that all passwords are crackable, so it is just a matter of time. If an attacker tries every possible combination, she will eventually guess the correct password. Usually with brute-force attacks, an attacker gains access to the encrypted passwords and downloads them to her local system. Then she tries every possible combination until she guesses the passwords. Remember, if an attacker has the encrypted passwords for every user on your system, she does not have to crack every password to get access—she only has to crack one.

A subset of the brute-force attack is the *dictionary attack*. If users have really strong passwords, attackers need to try every possible combination until they get access. But as was already discussed, users don't typically choose strong passwords. Most users pick very easy passwords based on dictionary words. In that case, instead of trying every possible combination, an attacker would try every word in a dictionary. This is a much smaller subset than every single possible combination, and because the attacker needs to crack only one or two passwords, her chance of success is very high.

Denial-of-Service

When most people think of attacks against a system, they think of someone trying to gain access. However, in some situations preventing others from gaining access can be just as useful. These types of attacks are *denial-of-service attacks*. If you are at a client's site, giving a demo to close a sale and you cannot get access to your system, that can be just as embarrassing and damaging as if your competition stole your material. There are several ways someone can launch a denial-of-service attack against access control. Most accounts are set up so that after a certain number of failed logon attempts, the account is locked. In this case, an attacker can just try to log on to every account, giving bad passwords, and lock every account on the system so no one can gain access. The other way is to flood the pipes so no one can even get access to the server.

This attack is popular with dial-up accounts. If an attacker knows that a company provides dial-up access, he keeps dialing the number and connecting to the modem pool from different computers. Eventually, he uses up all the phone lines, and legitimate users will be unable to gain access to the system.

Spoofing

When you were young and wanted to go to a club or bar with your friends, what did you do? You acquired a fake ID so you could pretend to be someone else who was older and could get access to a facility that you normally should not have access to.

When you acquired that fake ID, you were *spoofing* your identity. The same thing can be done with access control. An intruder would not normally be allowed access to your system. So, if he tried to authenticate as Joe Attacker, your system would deny him access. However, if he acquires the one-time password device for a given user and acts like that user (or *spoofs* that user), the system would give him access because the system thinks he is a legitimate user and does not know that he is really an attacker. This is the problem with access control that is based solely on something you have. If you have the device, the system will allow you in, but as you can see, it is very easy for someone else to acquire and gain access.

To prevent the spoofing attack, you should have multilevel access control. To gain access in such a system, you would need something you know and something you have. This way, even if an intruder can steal your device that generates your password, he will not be able to get access.

Sniffing

Some systems require that you have a user ID and a password to gain access, but they send the password over the network in plain text. An attacker can put a *sniffer* on the wire, which is a passive attack that allows her to watch the traffic going over the wire. Because the plain text is not encrypted, the sniffers can read the password and user ID and then use those passwords to gain access. It is very critical that any network authentication scheme encrypts the password before it sends the password over the wire.

MONITORING

Explain intrusion detection.

A key motto of security is “prevention is ideal, but detection is a must.” As long as you have a connection to an untrusted network like the Internet, you will not be able to block every attack. Some attacks will sneak in because you have to allow traffic to flow from a business standpoint. Even if you allow only port 80 traffic into a certain system, an attacker can still attack over that port, and your prevention measures (such as firewalls) will allow it through because they allow Web traffic to that given host. Therefore, you need someone or something to detect attacks in a timely manner. This is done by monitoring your systems and network traffic looking for unusual patterns or things that would be indicative of an attack.

Intrusion Detection

The field of study dealing with monitoring networks and hosts and looking for attacks is known as *intrusion detection*. The critical thing to remember with intrusion detection is that you are passively monitoring a network or hosts looking for signs of an attack.

The emphasis is on detection, not prevention. Inline devices such as firewalls actively block or allow traffic depending on various things like header information. Intrusion detection systems (IDSs) act more like sniffers in that, by themselves, they do not actually prevent attacks—they just alert that a potential problem exists. It is common practice to check the logs or set up alerts in an IDS for unauthorized access to certain resources. For example, most companies have a policy that no one should ever gain access as root to an internal resource from an external address. So, if your IDS ever sees someone remotely logging in as root, it should set off an alert that there is an access violation on the system.

Types of Intrusions

To better understand ways IDSs work, let's look at some of the types of intrusions and the impact they could have on your network. Intrusions can be categorized in many ways, but the following is one way of addressing the problem:

- ◆ Host versus network
- ◆ Passive versus active
- ◆ Known versus unknown

Host Versus Network

When an attacker tries to gain access to a company's infrastructure, the first question is what is she trying to attack. Is she trying to gain access to a specific host, or is she trying to gain access to an entire network? Depending on what she is going after, the intruder can be detected in several ways.

In addition to what the intruder is trying to attack, you can also ask what the means are by which the attack is going to be launched. In most cases when talking about the Internet, an attacker uses the network to launch the attack because it allows her to do it from anywhere. Attacking a company through a single host usually requires either gaining physical access to a facility or stealing a computer to gain access. These types of attack aren't less feasible; an attacker is just less likely to go through that much effort to gain access.

IDSs are typically broken down into host-based intrusion detection systems (HIDSs) and network-based intrusion detection systems (NIDSs). HIDSs sit on a single computer system and look for signs of an intrusion. They can usually be more finely tuned to a specific computer system but do not scale well across an enterprise that might have thousands of systems. NIDSs, on the other hand, sit on a network like a sniffer examining traffic for signs of an attack. They tend to look for more general types of attacks but scale very well because on one network segment they can protect thousands of systems.

Passive Versus Active

After an intruder gains access, what is he going to do? In most cases, an intruder is going to actively do something like deface a Web site, steal corporate secrets, or plant a back door on the system. In other cases, however, an attacker might just passively monitor traffic or keystrokes to try to gather information or a password for a particular account, such as root.

Active attacks are usually easier to detect because the intruder is actually doing something on the network. With passive attacks, the intruder is essentially just listening, and because he is not doing anything per se, detecting and stopping these types of attacks is much harder. The way to stop a passive attack is to not allow the intruder to get access in the first place.

Known Versus Unknown

This is a more abstract breakdown of intrusions, but it plays a key role in how easy detecting an attack is. A lot of known attacks are still used by intruders to break into systems. The best way to define a *known* attack is one in which the vendor has acknowledged a security hole in its software. Usually with known attacks, the vendor has also released a patch, so hopefully a company would apply the patch so that it would no longer be vulnerable. However, many known attacks still have a wide range of success because companies do not religiously apply patches to their systems. Known attacks are easy to detect because you know what the attack looks like.

Unknown attacks are attacks in which a small group of people know about the attack but it is not public knowledge, so the vendor does not know the vulnerability exists and therefore cannot release a patch. Unknown intrusions are very difficult to detect and prevent.

Intrusion Prevention

The term *intrusion prevention* has undergone changes in its meaning since early 2002. Prior to 2002, the main way to prevent an intrusion was to closely control access through strong identification and authentication. For example, instead of using weak passwords to gain remote access, you would use one-time passwords or biometrics, which are much harder for an attacker to defeat. Most of the emphasis has been on authentication because identification is usually through a password, which most companies make very predictable.

During the course of 2002, intrusion prevention has been used to describe a new class of systems that have grown out of the intrusion detection market. Firewalls are active devices in which traffic passes through. Usually based on header information, traffic can be either blocked or allowed. Intrusion detection systems were passive devices that would alert when an attack occurred but not actually stop the traffic. *Intrusion prevention* systems are a mixture of both. They work like a typical IDS looking for possible attacks on a network, but they are also active devices like firewalls through which traffic must pass. If the IDS senses an attack, instead of just alerting like it previously did, it can now actually stop the attack by blocking the traffic or preventing the malicious behavior by enforcing rules and policies.

How Intrusion Detection Works

Intrusion detection systems come in many shapes and sizes. At the most basic level is the question of where you should place the IDS. We have talked about network- versus host-based intrusion detection and the pros and cons of each. This section examines the two general types of IDS—signature matching and anomaly detection—focusing on how they operate and detect attacks.

Signature Matching

Signature or pattern matching IDS maintains a database of known attack signatures. When it looks at traffic for NIDS or at log files for HIDS, it tries to find a match for each of these signatures. If it finds a match, it sends off an alert that the system is being attacked. This approach is similar to how virus scanning software works. The virus software maintains a database of known viruses and looks for those patterns across all files.

The pattern matching approach has positive and negative aspects. The positive side is that they are fairly easy to update across a large number of companies. You essentially create a new signature and can push it out to every IDS. A company can also easily create its own signatures even if the IDS vendor does not present a signature. The negative side is that they detect only known attacks. If a new attack comes out, a signature-based IDS has no chance of detecting it. Another big drawback is that because they are based on static signatures, they tend to generate a high number of false alarms. This occurs when the IDS says it has found an attack, but in reality it is normal traffic.

Anomaly Detection

The concept behind anomaly detection is to determine what is normal traffic for a company and anything that falls outside that norm is deemed an attack and is dropped. The positive aspects of such an approach are very obvious. Because there are no signatures, you do not have to worry about constantly updating the system with new signatures. Also, because it is not based on signatures, it can detect both known and unknown attacks on a system. The disadvantage is figuring out how you should determine what is normal. Normal would be different for every company, and even within a company it is constantly changing over time. So, you need some way to learn the network for a given company and constantly change it over time.

Most systems are based on signature detection with some anomaly detection.

Now that we have looked at the main types of IDS, we will briefly cover how they operate. After an IDS determines that an attack has been detected, it sets off some type of alarm. Depending on the severity, this alarm can range from putting a message on a screen to sending an alert to someone's pager. Some IDSs can actually send messages to firewalls that will actually update their rule sets so that they can block these attacks in the future. Automatically changing a rule set on a firewall is very dangerous because it opens the door for an attacker to spoof an IDS and change the rule set. Even if the updates allow an IDS to only block traffic, an attacker could still launch a denial-of-service attack against a company by simulating an IDS and setting up rules to block traffic coming from anywhere.

PENETRATION TESTING

One error that companies often make is that they set up access controls and then test the access controls to make sure they are working properly. The problem with how companies approach this is that they usually test the positive but do not test the negative.

What I mean by that is that after they set up access controls, they test and make sure users can get to the resources they need to access. So, if Bob needs access to server A and C, they would test and see whether Bob could access both servers. If he could, they then conclude that the access controls have been set up properly. The problem with this is testing the negative—what else can Bob access? If Bob can also access server D, the company has given Bob too much access and is not adhering to a principle of least privilege. Testing the positive is easier because it is a smaller amount of testing. When you test the negative, you have to test all the possible combinations, which can be time-consuming. This is one reason companies do not do it.

The other reason they do not test the negative is if it is done incorrectly, there is less of a pain factor. For example, if Bob is supposed to have access to server A and does not have access, he might get upset, complain, and cause a lot of problems. The company will also get upset because they are paying Bob to do a job, and you are preventing him from doing that job. If Bob is not supposed to have access to server D and he does, there is a good chance he will not even notice—and if he notices, he will probably not tell anyone.

Instead of trying to manually test the negative, a good alternative is to perform penetration testing to check the access controls of a company. Penetration testing is sometimes referred to as *ethical hacking* because you are trying to simulate how an attacker would break into a system and find the holes before an attacker does. The idea is that if you try to break into your own systems, you can find weaknesses in your access control policy and fix them before a real attacker breaks in.

Penetration Testing Versus Security Assessments

Penetration testing is sometimes contrasted or compared with *security assessments*. The main difference between the two has to do with the scope and amount of initial information one is given. Typically, with a penetration test (or pen test), you are testing the security from the Internet so you are given a domain name and maybe an IP address but nothing else. The goal is to see how much you can find out about the company, including possible ways you can break in. The problem is that some companies think that if you are able to get access to the system, you have proven that the security is weak. The main issue with this is that you do not get a complete picture of your overall security. You know that there is one way into your system, but are there others? It does not give you a comprehensive view of the current weaknesses in security across your company. The second issue is that it does not include threat or risk assessment, which play critical roles in a company's security.

Security assessments usually include a penetration test but are much more thorough. You are typically given access to all the key systems within a company to evaluate the current level of security. With security assessments, you are not trying to prove that you can get in; you are trying to paint a picture of the current threats that exist to the organization and what needs to be done to protect against them.

Ethical Issues

Whenever you talk about breaking into a system, there are always ethical issues surrounding this. Is it ethical to try to break into a company even if you are not going to do harm? Is it ethical to probe a system even if you do not have permissions? The first rule of thumb is to always get permission in writing before you do any form of penetration testing. Before you even think about doing anything against a company, you always need to get permission in writing. A point that some people bring up is that a company does not always want to know when you are doing a penetration test. If the people responsible for security know on a given day at a given hour someone is going to try to break in, there is a good chance they might temporarily increase their security to skew the test results. Even if this is what a company wants to do, you can still get permission in writing.

It can say that over the next five days, this will be performed, or it can be signed by the CTO who decides not to tell his staff this is being done.

I keep emphasizing the “in writing” part of this discussion. Verbal contracts are binding IF you can prove them. But proving you have a verbal agreement to do something is very difficult. If you have a signed piece of paper, the opposing party will have a hard time denying that they agreed to something. Always err on the side of caution and get permission in writing.

Some people argue that if the systems are connected to the Internet and you are only going to probe and not do any damage, you don't need to get permission. This tends to be a big ethical issue at security conferences, but to me the answer is simple: The system does not belong to you. It is someone else's system, and if you want to do something to it, you need to get the owner's permission. The other problem is that, because you are remotely probing a system, you might not intend to do any damage but by accident might crash or reboot the system. If you get the company's permission, it might have you perform a penetration test during low-volume hours so if something happens, the financial impact is minimal. If you just decide to probe a company without permission, you could crash the system and cause a large financial loss to a company.

Performing a Penetration Test

The best way to simulate an attack is to follow the same process that an attacker takes to break into a system.

As noted in the popular *Hackers Beware*, the following process, outlined in Step By Step 1.1, is used by attackers to break into systems.

STEP BY STEP

1.1 The Process for Breaking into a System

1. Perform passive reconnaissance.
 2. Perform active reconnaissance (scanning).
-

3. Exploit the system by gaining access through the following attacks:
 - Operating system attacks
 - Application-level attacks
 - Scripts and sample program attacks
 - Misconfiguration attacks
 - Elevating of privileges
 - Denial-of-service attacks

 4. Upload programs.

 5. Download data.

 6. Maintain access by using the following:
 - Back doors
 - Trojan horses

 7. Cover your tracks.
-

In most cases when performing a penetration test, to check security and access control you would perform only steps 1–3. After you prove you can get into the system, you would stop. In other cases, companies want you to see whether they can find you and how much data you can get, so you would continue with steps 4–7.

Common Tools

Several tools are available that can be used to perform penetration testing. This section briefly covers two common tools. The tools you choose depend on how manual or automated you want to make the process. Manual processes involve taking the output of each tool and manually probing into each port looking for potential areas to exploit. This approach takes longer but has a higher chance of finding more vulnerabilities. The other approach is more automated and involves using vulnerability scanners that scan a given set of addresses looking for known vulnerabilities.

NOTE

Nessus One of the most common free vulnerability scanners is Nessus. Nessus scans for several hundred vulnerabilities across various operating systems and reports back on which vulnerabilities are open on a given system. The key thing to remember with vulnerability scanners is that they detect only known vulnerabilities.

One of the most common port scanners is nmap. nmap not only determines which ports are open, but also performs OS fingerprinting and other advanced features such as sending out decoy packets to spoof who the real attacker is.

CASE STUDY: THE SMART CARD CASE

ESSENCE OF THE CASE

The following are the essence of the case:

- ▶ Eventually it was determined that intruders had obtained a certificate that enabled them to install their own certificate authority (CA) and produce smart cards trusted by the ABC Company's computer systems.
- ▶ Among other capabilities, the CA is the computer in the public key infrastructure (PKI) that issues certificates. In the ABC PKI, the certificates are used on smart cards, and in this particular PKI implementation, a hierarchical structure is allowed. In other words, the root—or first CA—can produce a certificate that authorizes another CA. Smart cards produced by either CA can then authorize access to computer systems.
- ▶ The intruders were able to obtain a certificate from the first CA, install their own CA, and produce smart cards that they then used on the system.

SCENARIO

ABC Company (not a real name) recently instituted a smart card program. All employees are required to use smart cards for access to data systems. Authentication and identification information is placed on the smart card and used to log the user onto the computer. A smart card and associated PIN number are necessary for logon.

Smart cards are issued by the human resources department when an employee is hired and can be reissued as required. The cards are also used for access to the building and contain a picture of the employee. Cards must be controlled by the employees at all times. If an employee leaves his desk, he must remove the card from the smart card reader and carry it with him for identification. Removing the smart card locks the computer and prevents unauthorized users or intruders from accessing systems when an employee is away from his desk. This also prevents smart card sharing because the card must be in the reader for the computer session to remain accessible to the user.

This excellent system of access control has many features that make it desirable. The automatic logoff, identification requirements, building access requirements, and one-user-one-device requirement all make it an outstanding design.

CASE STUDY: THE SMART CARD CASE

Unfortunately, a routine audit has disclosed that multiple logons by the company's vice president were made when she was on vacation. She was able to prove that her smart card was in her possession at the time the intruder was using a smart card issued to the VP's account to access the network. Further research uncovered the use of multiple "fake" smart cards to access the accounts of other privileged users and thus provide access to other sensitive documents.

ANALYSIS

This is an example of why exotic and complicated technical systems are not the end all and be all of security. In this case, the root was not appropriately protected. Even though PKI can provide a strong authentication and access control system, it is reliant on human beings to design a secure PKI.

CHAPTER SUMMARY

Access control compliments other areas of security but is critical to achieving defense in depth across your organization. Without access controls, you are saying that after someone gets access to a system, he can do whatever he wants because there is nothing restricting his actions. This chapter outlined various approaches to access control and how it can be achieved across an organization.

KEY TERMS

- Access controls
- ACLs
- Bell-LaPadula
- Biba
- Brute-force attack
- Denial-of-service attack
- Discretionary access control (DAC)
- IDS
- Lattice-based access control
- Liptner
- Mandatory access control (MAC)

CHAPTER SUMMARY

KEY TERMS

- Nessus
- Non-inference
- Penetration testing
- Role-based access control
- Rule-based access control
- Signature matching
- Sniffing
- Spoofing
- SSO
- Star property
- Trusts

APPLY YOUR KNOWLEDGE**Exercises****1.1 Rule-Based or Role-Based: Which Is It?**

Examine the access control system of a Windows NT or Windows 2000 system. Determine whether it is role-based or rule-based, and explain why.

Estimated Time: 20 minutes

1. Examine the default user groups on the system. What groups exist? Do they have specific rights or access that is allowed on the system?
2. Determine whether additional groups can be created. Who can create these groups? Can rights or access be granted to these groups?
3. Determine whether individual user accounts can be given rights and access on the system.
4. Based on your study, is this a rule-based or role-based system of access control? Why?

Answers to the exercise:

1. Multiple user groups exist, depending on whether you are looking at Windows NT or Windows 2000 and whether the computer is a domain controller, server, or workstation. All domain controllers have the Administrators, Account Operators, Server Operators, Print Operators, Backup Operators, Domain Guests, and Domain Users groups. Each default group has specific rights assigned to it, and access control lists on resources determine which groups have which type of access.
2. Additional groups can be created and granted rights and access to resources.
3. Individual user accounts can be given rights and access on the system, either on their own or by membership in groups.

4. This can be a rule-based or a role-based system of access control depending on its implementation. Clearly, default groups are granted access dependent on presumed roles. Additional groups can also be assigned roles and granted associated rights and access. However, there is no enforcement of these roles because enforcement is based on human interaction. If the policy is strict and followed faithfully, a user is given access according to the role he plays by his inclusion in a group that has only the access and rights he requires to perform his functions. Rule-based control can also be implemented by writing rules for each user's access and implementing it by assigning his individual account the right or access outlined in the rule developed to govern his behavior on the system.

Review Questions

1. What is the correct policy to use for shared accounts?
2. Describe the difference between discretionary access controls and mandatory access controls.
3. Lattice-based access control is a form of MAC. Flow operations for this type of MAC include the properties of partial order, which are what?
4. Collections of rules that apply to network access through a router based on IP address or port are _____.
5. The first time someone logs onto a new account, she should be forced to change her password. This is for what reason?

APPLY YOUR KNOWLEDGE

6. The information access model that is meant to protect against write-down Trojan horses is the _____ model. In this model a user with high privileges will not be able to write to areas where only a lower privilege is necessary.
 7. Explain the difference between identification and authentication.
 8. What problems do one-time passwords solve?
 9. What is one problem with single sign on?
 10. What is the usefulness of TACACS+ and RADIUS?
 11. Explain how a brute-force attack can be used to crack passwords.
 12. Define intrusion detection and give an example of where it is useful.
 13. What is the difference between host and network forms of intrusion detection?
3. Which principle makes people respond to access controls?
 - A. Accountability
 - B. Authentication
 - C. Authorization
 - D. Accreditation
 4. A user can have multiple levels of access to a system depending on the work that she must do. In a MAC system, this might mean that she could log on at her highest level of access to do all her work. What can be done to correct this limitation of MAC controls?
 - A. Never give a user more than one level of access control.
 - B. Audit the use of her access and punish her for using her higher level access logon when it is not necessary.
 - C. Use an access level system (compartmentalization) that is not all inclusive—that is, a higher-level access account cannot access lower-level resources.
 - D. Only give her the highest level access logon she needs. She can access anything she needs to access with this. Why give her multiple accounts?

Exam Questions

1. Which principle identifies a user and verifies that the user is who he says he is?
 - A. Authentication
 - B. Access control
 - C. Biba
 - D. Bell-LaPadula
2. Which principle determines what resources the user can use on the network?
 - A. Authentication
 - B. Access control

APPLY YOUR KNOWLEDGE

- B. Rule-based access control is necessary for small businesses, whereas role-based access control is necessary for large businesses.
- C. Rule-based access controls assign access parameters to user accounts, whereas role-based access control is based on access control desired according to the job function of a position.
- D. Rule-based access controls are easy to implement, whereas role-based access controls are not.
6. When assigning access to sensitive information you should maintain which of the following?
- A. Separation of duties
 - B. One account, one user
 - C. Least privilege
 - D. Accountability
7. When assigning permissions to accounts, you should give the access that the user needs and nothing more. This defines which security principle?
- A. Separation of duties
 - B. One account, one user
 - C. Least privilege
 - D. Accountability
8. The access control model that defines simple security as the reading of files and the star property with writing of files is which of the following?
- A. Biba
 - B. Bell-LaPadula
 - C. Liptner
 - D. Non-inference
9. Which model deals with integrity instead of confidentiality?
- A. Biba
 - B. Bell-LaPadula
 - C. Liptner
 - D. Non-inference
10. Which model applies government models to commercial settings?
- A. Biba
 - B. Bell-LaPadula
 - C. Liptner
 - D. Non-inference
11. Which access control model deals with the information you can find out by observing the input to and output from a system?
- A. Biba
 - B. Bell-LaPadula
 - C. Liptner
 - D. Non-inference

Answers to Review Questions

1. Account sharing is not allowed. When accounts are shared, there is no accountability. For more information, see the “Accountability” section.

APPLY YOUR KNOWLEDGE

2. Discretionary access controls are based on human decisions. Policy determines whether a user, a service, or an application can access a resource such as a file or directory. It does not provide a high level of access control because the measure of who should have access is subjective—a human gives and takes controls. Mandatory access controls are done at a higher level: The computer system is in control. Entities that use the system are given a classification level which is associated with their accounts. Data also has a classification level. The system determines access by looking at the classification of the user and the data. For more information, see the “Discretionary Access Control” and “Mandatory Access Control” sections.
3. Reflexive, antisymmetric, and transitive. For more information see the “Lattice-Based Access Control” section.
4. Access control lists. For more information, see the “Access Control Lists” section.
5. The default password used to log on might be known to others. The use of authentication and identification to control access works only if the individual who owns the account is the only one who knows its password. This also enables accountability. For more information, see the “Account Administration” section.
6. Bell-LaPadula model. For more information, see the “Access Control Models” section.
7. Identification is the presentation of credentials that identify who the user is. The user account ID is an identification credential. Authentication is the process of proving that the user is who he says he is, often by using a password or other piece of information known only to this user.

For more information, see the “Identification and Authentication Techniques” section.
8. One-time passwords solve the problem of weak passwords, or shared passwords. When passwords are used, they are good only if they are known only to the user. Often users write down passwords or share them. Passwords can also be cracked by programs built to do so. One-time passwords are only good when used, thus it doesn’t matter if they’re captured or written down because they cannot be reused. For more information, see the “Identification and Authentication Techniques” section.
9. Single sign-on means that one user ID and password provide access to all the network resources assigned. Unfortunately, it also means that one compromise of that network ID and password means the intruder has acquired access to all the resources assigned. For more information, see the “Single Sign-on” section.
10. TACACS+ and RADIUS provide centralized authentication. This can be used to provide authentication to multiple applications or to the network from remote access. For more information, see the “Centralized/Remote Authentication Access Controls” section.
11. A brute-force attack is one that tries all possible combinations to determine a password. Password crackers often operate in this mode, trying every possible character combination until the password is matched. For more information, see the “Brute-Force” section.
12. Intrusion detection is the capability to detect when unauthorized access is taking place or has taken place. This is useful because it can identify an attack in progress, in which case, perhaps the attacker’s success can be limited or his information can be gathered for later prosecution.

APPLY YOUR KNOWLEDGE

It is also useful because it can indicate what the attacker accessed and what information he obtained. For more information, see the “Intrusion Detection” section.

13. Host intrusion detection places agents on the host machine and records when the host has been accessed in an unauthorized manner. Network intrusion detection agents listen to all network activity and can find when any intruders have accessed the network. For more information, see the “Intrusion Detection” section.

Answers to Exam Questions

1. **A.** Answer B, access control, is the ability to control who and what resources are accessed. Answers C and D are incorrect because they are particular access control methodologies. See the “Introduction” section for more information.
2. **B.** Answer A is the process of proving you are who you say you are, so it’s wrong. Answers C and D are specific access control models, so they are incorrect.
3. **A.** Answer B is the process of proving you are who you say you are, so it’s incorrect. Answer C is the process of seeing if you should get access, so it’s incorrect. Answer D is incorrect because accreditation is the approval of specific criteria as developed by an accrediting agency. See the “Accountability” section for more information.
4. **C.** It might be impossible to never give a user more than one level of access control, so answer B is incorrect. Likewise, answer A might help but will not prevent the access, so it’s incorrect.
5. **C.** Answer A is incorrect because rule-based access control more often applies to users instead of groups. Answer B is incorrect because even small businesses might find rule-based access control difficult to manage, and answer D is incorrect because rule-based access controls can be difficult to implement when more than a few users are present. See the “Rule-Based Access Control” and “Role-Based Access Control” sections for more information.
6. **A.** Answers B, C, and D are incorrect because they are true for access control for all users, not just those of sensitive information. See the “Account Administration” section for more information.
7. **C.** Answer A is incorrect because it keeps a user from taking advantage of his access to sensitive information—the one who writes the code does not get to configure the system, and the one who approves the purchase of vendor goods does not get to issue the checks. Answer B means that accounts should not be shared, so it’s incorrect. Answer D provides control over the use of resources—if you access a resource, that access can be recorded—so it’s wrong. See the “Account Administration” section for more information.
8. **B.** All other models do not have this property, so answers A, C, and D are incorrect. See the section “Access Control Models” for more information.

Answer D is wrong because you should not be cavalier about this access—when a privileged user accesses an area of less privilege, she can infect the area of less privilege. See the “Mandatory Access Control” section for more information.

APPLY YOUR KNOWLEDGE

9. **A.** All other models do not have this property, so answers B, C, and D are incorrect. See the section “Access Control Models” for more information.
10. **C.** Answers A and B represent government access control models, so they’re wrong. Answer D represents a generic access control model, so it’s also wrong. See the section “Access Control Models” for more information.
11. **D.** All other models do not have this characteristic, so answers A, B, and C are incorrect. See the “Access Control Models” section for more information.

Suggested Readings and Resources

- Black, David K. “Confounding Access,” *infosecuritymag.com*, April 2002.
- Chauhan, Abishek. “Do Firewalls and IDS Create a False Sense of Internal Security?” *scmagazine.com*, September 2002.
- Hey, Wilf. “Securikey Authentication System,” *scmagazine.com*, June 2002.
- Kurzban, Stanley. “Implementation of Access Controls.” In *Handbook of Information Security Management*, edited by Micki Krause and Harold Tipton, Auerbach, 1999.
- Richards, Donald R. “Biometric Identification.” In *Handbook of Information Security Management*, edited by Micki Krause and Harold Tipton, Auerbach, 1999.
- Ross, Leo. “Single Sign-on.” In *Handbook of Information Security Management, Fourth Edition, Volume 2*, edited by Micki Krause and Harold Tipton, Auerbach, 2001.
- Smith, Richard. “The Strong Password Dilemma,” *Computer Security Journal*, Volume XVIII, Number 2, Spring 2002.
- Stackpole, Bill. “Centralized Authentication Services (RADIUS, TACACS, DIAMETER).” In *Handbook of Information Security Management, Fourth Edition, Volume 2*, edited by Micki Krause and Harold Tipton, Auerbach, 2001.
- Vallabhanein, S. Rao. “Access Control Systems and Methodology.” In *CISSP Examination Textbooks, Volume 1*, SRV Publications, 2000.
- <http://www.acm.org/sigsac/> (information on the ACM Special Interest Group on Security, Audit, and Control [SIGSAC]).
- [http://www.list.gmu.edu/journals/computer/pdf_ver/i931bacm\(org\).pdf](http://www.list.gmu.edu/journals/computer/pdf_ver/i931bacm(org).pdf) (“Lattice Based Access Control Models,” an article by Ravi S. Sandu).
- <http://www.microsoft.com/windowsxp/pro/using/howto/security/accesscontrol.asp> (“Use Access Control to Restrict Who Can Access Files,” an article on XP file access control).

OBJECTIVES

Identify the key areas of knowledge of telecommunications and network security.

Explain the International Standards Organization/Open Systems Interconnection (ISO/OSI) layers and characteristics including:

- **Physical layer**
 - **Data Link layer**
 - **Network layer**
 - **Transport layer**
 - **Session layer**
 - **Presentation layer**
 - **Application layer**
- The ISO/OSI seven-layer model defines the fundamental aspects of how all network communication occurs. The seven layers are presented as a framework that networking vendors use to ensure interoperability between platforms and protocols. Understanding how network communications is defined allows the security professional to understand where the implications of security exploits may occur.

Describe the design and function of communications and network security including the following:

- **Physical media characteristics (for example, fiber optics/coaxial/twisted pair)**
- **Network topologies (for example, star, bus, and ring)**
- **IPSec authentication and confidentiality**



CHAPTER 2

Telecommunications and Network Security

OBJECTIVES

- **TCP/IP characteristics and vulnerabilities**
 - **Local area networks (LANs)**
 - **Wide area networks (WANs)**
 - **Remote access/telecommuting techniques**
 - **Secure Remote Procedure Call (S-RPC)**
 - **Remote Access Dial-In User System/
Terminal Access Control**
 - **Access system (RADIUS/TACACS)**
 - **Network monitors and packet sniffers**
- To properly secure networking communications, you must understand how networks are designed and how communications occur across networks. By understanding the design principles and functions of different networking technologies, the security professional can better understand how to properly secure those technologies.

Describe the components, protocols and services involved in Internet/intranet/extranet design including the following:

- **Firewalls**
- **Routers**
- **Switches**
- **Gateways**
- **Proxies**
- **Protocols**
 - **Transmission Control Protocol/Internet Protocol (TCP/IP)**
 - **Network layer security protocols (IPSec, SKIP, SWIPE)**

- **Transport layer security protocols (SSL)**
- **Application layer security protocols (S/MIME, SSL, SET, PEM)**
- **Challenge Handshake Authentication Protocol (CHAP)**
- **Password Authentication Protocol (PAP)**
- **Point-to-Point Protocol (PPP)/Serial Line Internet Protocol (SLIP)**
- **Services**
 - **High-level Data Link Control (HDLC)**
 - **Frame relay**
 - **Synchronous Data Link Control (SDLC)**
 - **Integrated Services Digital Network (ISDN)**
 - **X.25**

- After a security professional understands the network design concepts, she must then understand the components, protocols, and services that enable the communications to occur. The methods of securing a router are not necessarily the same as securing a switch. Knowing this enables the security professional to select the proper methods of securing her network components, protocols, and services.

Define and describe communications security techniques to prevent, detect, and correct errors so that integrity, availability, and confidentiality of transactions over networks may be maintained:

- **Tunneling**
- **Virtual Private Network (VPN)**
- **Network monitors and packet sniffers**

OBJECTIVES

- **Network Address Translation**
 - **Transparency**
 - **Hash totals**
 - **Record sequence checking**
 - **Transmission logging**
 - **Transmission error correction**
 - **Retransmission controls**
- Today's complex networks almost require security professionals to operate their networks in conditions that are less than ideal security conditions. To address this, there are a number of methods of mitigating the risk of the requirement of exposing network resources. Understanding how to implement designs such as tunnels and VPNs, as well as knowing how to determine if the traffic is indeed protected, helps to ensure that the security level of traffic and transactions in "hostile" environments is protected.

Define and describe specific areas of communication and how they can be secured:

- **Email security**
 - **Facsimile security**
 - **Secure Voice Communications**
 - **Security boundaries and how to translate security policy to security controls and practical application**
- Certain types of communications must occur between remote destinations. The problem with this is that it is difficult to ensure the security of these communications methods because they typically traverse insecure network links and segments.

In addition, because the communications will often come from a remote location, there is the risk of how to safely enable the communications to occur. Understanding how communications methods like email, facsimile, and voice communication occur will help the security professional understand how to secure this traffic.

Explain current forms of network attacks and their countermeasures including

- **Address Resolution Protocol (ARP)**
 - **Brute force**
 - **Worms**
 - **Flooding**
 - **Eavesdropping**
 - **Sniffers**
 - **Spamming**
 - **Private Branch Exchange (PBX) Fraud and Abuse**
- There is an old saying, "Know thine enemy." This holds true in securing telecommunications and network security. Security professionals do not need to be "hackers," but understanding the nature of different types of network attacks and exploits will assist in a security professional's ability to recognize and protect against such attacks.

OUTLINE

Introduction	67	Mesh Topology	93
		LAN and WAN Technologies	94
		Ethernet	95
		Token-Ring and FDDI	98
		Attached Resource Computer Network	99
The Open Systems Interconnection Model	68	LAN Devices	99
The OSI Layers	70	Hubs and Repeaters	99
Application Layer	72	Switches and Bridges	100
Presentation Layer	73	VLANs	101
Session Layer	73	Routers	103
Transport Layer	74	Firewalls	104
Network Layer	75	Gateways and Proxies	110
Data Link Layer	75		
Physical Layer	76	WAN Technologies	110
OSI Summary	77	Dedicated Connections	111
Network Characteristics and Topologies	78	Circuit-Switched Connections	113
Coax	79	Packet-Switched Connections	113
10BASE-2 Specifications	80	Cell-Switched Connections	114
10BASE-5 Specifications	81	WAN Services	114
Unshielded Twisted Pair	82	Point-to-Point Protocol and Serial	
UTP Specifications	83	Line Internet Protocol	114
Troubleshooting UTP	84	High-Level Data-Link Control	115
Fiber Optic	84	X.25	115
Fiber-Optic Cable Components	84	Link Access Procedure Balanced	116
Multi-Mode Fiber	86	Frame Relay	116
Single-Mode Fiber	86	Synchronous Data-Link Control	116
Dense Wave Division Multiplexing	87	Integrated Services Data Network	116
Wireless	87	Digital Subscriber Line	117
		Switched Multimegabit Data Service	118
		High Speed Serial Interface	118
Network Topologies	89	WAN Devices	118
Linear Bus Topology	89		
Star Topology	91		
Ring Topology	92		
Tree Topology	93		

OUTLINE

Providing Remote Access Capabilities	119	Network Monitoring and Packet Sniffers	137
Client-Based Dial-in Remote Access	119	Intrusion Detection	139
Using Tunneling As a Security Method	120	Intrusion Response	141
Virtual Private Networks	121	Network Address Translation	142
Client-Based VPNs	121	Transparency	144
Site-to-Site VPNs	122	Hash Totals	145
VPN Protocols	123	Email Security	146
Remote Access Authentication	124	Facsimile and Printer Security	147
		Common Attacks and Countermeasures	147
Networking Protocols	125	Class A Abuses	147
Transmission Control Protocol/Internet Protocol	125	Class B Abuses	148
Application Layer Protocols	126	Class C Abuses	149
Transport Layer Protocols	127	Class D Abuses	150
Reviewing TCP and UDP	129	Class E Abuses	152
Internet Layer Protocols	129	Class F Abuses	154
		Fault Tolerance and Data Restoration	155
Protecting the Integrity, Availability, and Confidentiality of Network Data	130	Managing Network Single Points of Failure	158
The CIA Triad	130	Cable Failures	158
Security Boundaries and Translating Security Policy to Controls	132	Topology Failures	159
Trusted Network Interpretation	133		
Network Layer Security Protocols	135	Chapter Summary	163
Transport Layer Security Protocols	136		
Application Layer Security Protocols	136	Apply Your Knowledge	165

STUDY STRATEGIES

- ▶ The Telecommunications and Network Security domain is a positively massive amount of data to cover. Ranging from the structure of networking frameworks, to network topologies, to network devices to security practices, there is a wide playing field to cover.
- ▶ The best way to approach the subject is to focus on the individual sections instead of trying to understand the entire domain at one time. Break the domain into logical groupings of topics. I like to start with the OSI model because it provides the foundation for networking in the first place.
- ▶ Use the layered approach of the OSI model to focus on the specific technologies and concepts. Start with layer-1 concepts like network cabling and physical design. Move up to network functions at layer 2. Proceed to layer-3 concepts, and so on.
- ▶ Try to focus your LAN and WAN study topics. Work on mastering the various LAN devices and technologies, and then proceed to the WAN devices and technologies.
- ▶ After you lay the foundation of understanding the fundamental networking concepts, proceed to the more complex security discussions. Start easy and look at the security theory and practices before you proceed to the more specific security threats and countermeasures.
- ▶ Above all else though, remember to take small steps. “Grasshopper, first you must take the stone, and then you can go.” Keep this philosophy in mind. Master a concept before you attempt to proceed to the next one.

“Telecommunications and Network Security domain encompasses the structures, transmission methods, transport formats, and security measures used to provide integrity, availability, authentication, and confidentiality for transmissions over private and public communications networks and media.

The candidate is expected to demonstrate an understanding of communications and network security as it relates to voice communications; data communications in terms of local area, wide area, and remote access; Internet/Intranet/Extranet in terms of Firewalls, Routers, and TCP/IP; and communications security management and techniques in terms of preventive, detective and corrective measures.”

—Common Book of Knowledge study guide

INTRODUCTION

This chapter explores the devices and technologies that constitute and define networks. We start with an examination of the Open Systems Interconnection (OSI) model and how it facilitates network communications. We then look at the network characteristics and topologies, including local area network and wide area network devices, services, and protocols. We will also define what a firewall is and is not, and look at methods of providing remote access to internal resources. After we have defined the things that constitute a network, we will start looking at methods of protecting the data and resources that run on our networks. We will finish with a look at fault tolerance and data redundancy.

As mentioned, the Telecommunications and Network Security domain is a very broad topic to discuss. This chapter has been broken down into numerous sections to make it easier to understand all the components of this domain and how they fit together.

THE OPEN SYSTEMS INTERCONNECTION MODEL

The early need for network computers was born out of the desire to share resources, specifically, printer resources. During the mid-1980s to early 1990s, very few systems were networked. This was due in large part to incompatible technologies. Companies started to recognize that they needed to buy a printer for each employee, even though each employee typically used the printer infrequently. Simply put, it was bad business. Companies decided it would make sense to share the printer among multiple users, thereby reducing costs and overhead. The early corporate networks were largely glorified methods to share printer resources.

As time progressed though, companies started to consider sharing other resources. They found that many times people would type a document, print it, and then give it to someone who would type it back into their system. Companies figured if they could share printers, why couldn't they share data? At that point the days of sneaker-net began their rapid demise leaving us with what we now take for granted: instant access to global resources.

It wasn't easy going though. One of the biggest hindrances to networking was the lack of standards. Everyone had a different method to network, and none of them worked well (if at all) with each other. This wasn't limited to just topologies, though. Some network interface cards (NICs) would only run Internet protocol (IP) or internetwork packet exchange (IPX) because of driver limitations. There was no ability to run both IP and IPX at the same time, so a decision had to be made as to what protocol all the systems would use. Likewise, clients could only connect to Novell or Unix or Microsoft at a time. As a result you needed to decide what everyone would use. Either everyone connected to Novell servers or everyone connected to Microsoft servers, and so on. This created a *monolithic networking model* that did not scale at all. To address these issues, various networking groups came together to create a scalable open standard that would facilitate the open communications between all systems. This became known as the OSI model.

Chapter 2 TELECOMMUNICATIONS AND NETWORK SECURITY 69

As I mentioned, in the early days of networking, systems were incompatible with each other. If you ran an IBM solution, you couldn't run a DecNet solution and vice versa. As a result, in the late 1970s the Open Systems Interconnection (OSI) model was created by the ISO to remove the barriers that hampered interoperability of network devices. Although the OSI model was a great idea, it is now more than 20 years old, and it is still a work in progress. Like they say, the great thing about standards is that there are so many to choose from!

The most difficult part of understanding the OSI model is recognizing that it is a framework of how networking functions, not a literal definition of how networking occurs. There is not necessarily a one-to-one mapping of layers to protocols. The OSI model exists to allow the user to understand the totality of a very complex system of communications by breaking the overall transmission of data into seven easier-to-define layers.

The easiest way I have found to understand and apply the OSI model is to do what I call "thinking layered," or my "elephant approach" to networking. Let's say you decide one day that you want to have elephant for dinner. If you decide to sit down and tackle an entire elephant all at once, you probably are not going to get very far. However, if instead of trying to do it all you sit down with nice easy-to-digest elephant steaks, before you know it you have the whole elephant taken care of. Applying the OSI model works in the same way. Rather than trying to understand the totality of network communications, try to break apart the communications into their layers ("think layered") and focus on understanding how each component works. Before you know it, you will have the complex network communications functions nailed down.

The OSI model really becomes clear if you run a network sniffer, decode the packets, and try to understand how what you are seeing applies to the OSI model.

The OSI model has become the primary model for architecting network systems. Rather than defining what should be done to facilitate network communications, OSI simply sets the expectations of what systems should expect to occur. It describes how data and network information should be communicated from the applications on one system to the applications on another system without stating what should be done to accomplish this.

The OSI Layers

The OSI reference model breaks this network methodology into seven separate layers. First you must understand that a reference model is simply a logical blueprint on how communications should take place. To address the processes that are required to communicate, OSI breaks the processes into logical groupings referred to as layers. These layers specify that each layer should be responsible for its own tasks and be able to interface with the layers directly above and below.

In a sense the layers are like departments in a company. A large software company has many different departments that facilitate the release of a product and the generation of revenue. The company wants to sell a product to a customer and the customer wants to buy a product that meets some defined need. The marketing department is responsible for determining what the customer needs are and presenting a marketing requirements specification to product development. Product marketing does not care how the customer need is met, as long as it is met. Product development is responsible for figuring out how to build and design a product that meets the customer requirements. Development is not interested in the conversations that take place between marketing and the customer, nor are they concerned with how the product will be sold. After the product is ready, the sales department works on determining the sales strategy and competitive analysis of the product. Sales doesn't care how the product was written, or in many cases what it does; instead sales focuses on how to sell the product to the customer. Each department has its own particular focus and function. On its own, each department is effectively worthless. Combined, however, the departments complement each other in delivering a total solution. A layered reference model is similar in concept. Some of the biggest benefits of a layered reference model are

- ◆ It divides the complex network operation into smaller, easier-to-manage pieces or layers; In our example, it is easier to manage the individual groups (marketing, sales, development, and so on) than to try to manage them all as a single thing.
- ◆ It facilitates the ability to make a change at one layer without having to change all the layers. This facilitates the ability to specialize the design and development of applications and protocols to specific tasks. In our example, the sales group can change its sales strategy without affecting how any other group performs its job.

- ◆ Defines a standard interface for multi-vendor integration. By using a standard interface, the details of how a particular layer functions are hidden from all the other layers, thus being transparent and allowing for multiple applications or protocols to function in concurrence; in our example, marketing can do whatever it wants to get the information it needs. Only as long as it always presents a marketing requirements specification, however, development will always know how to deal with the information it is presented.

It's important to understand that OSI does not define how to perform requisite tasks at each layer. This responsibility is left up to the individual vendors and the respective protocols. The OSI model simply defines what the expectations of each layer are, leaving the vendors and protocols to determine the best way to meet that expectation. As discussed, the OSI model is separated into seven distinct layers, as shown in Figure 2.1. Each layer has a core set of tasks and functions that it is responsible for providing. These layers are as follows:

- ◆ **Application layer (Layer 7)**—Primarily responsible for interfacing with the user. This is the application interface that the user experiences.
- ◆ **Presentation layer (Layer 6)**—Primarily responsible for translating the data from something the user expects to something the network expects.
- ◆ **Session layer (Layer 5)**—Primarily responsible for dialog control between systems and applications.
- ◆ **Transport layer (Layer 4)**—Primarily responsible for handling end-to-end data transport services.
- ◆ **Network layer (Layer 3)**—Primarily responsible for logical addressing.
- ◆ **Data Link layer (Layer 2)**—Primarily responsible for physical addressing.
- ◆ **Physical layer (Layer 1)**—Primarily responsible for physical delivery and specifications.




FIGURE 2.1
The OSI model.

Although there are seven distinct layers, it is important to understand that it does not necessarily mean that seven different protocols or applications are in use. Sometimes a single protocol may perform multiple functions across multiple layers. Remember, this is an architectural model not a literal model.

Let's look at each layer's function in more detail.

Application Layer

The Application layer is primarily responsible for providing the user access to network resources via the use of network-aware applications. The Application layer handles identifying and establishing that network resources are available. It is important to note that not every application—for example, word processing applications—is defined at the Application layer. Word processors do not have native networking functions, and thus are not network aware. On the other hand, World Wide Web (WWW) applications—for example, Web browsers—are network aware and thus are defined as Application layer entities. Some other examples of Application layer entities are

- ◆ **Email gateways**—Using Post Office Protocol (POP3), Simple Mail Transfer Protocol (SMTP), or X.400, email gateways deliver messages between applications.
- ◆ **Newsgroup and Internet Relay Chat (IRC) programs**—Using Network News Transfer Protocol (NNTP) and IRC, these applications provide for communications between hosts by allowing for either the posting of messages to a news server or the typing of a live conversation between chat clients.
- ◆ **Database applications**—Providing data storage and warehousing capabilities in central data repositories that can be accessed, managed, and updated.
- ◆ **WWW applications**—Providing access to Web resources, WWW applications include client Web browsers and Web servers.

Presentation Layer

The Presentation layer is often referred to as the “translator” of the network, similar to EBCDIC (Extended Binary-Coded Decimal Interchange Mode) and ASCII (American Standard Code for Information Interchange). As the name would imply, the primary purpose of the Presentation layer is to take data that is in a format the user understands and translate it into something that the network understands, and vice versa. In other words, it is presenting the data in the format that the next layer needs. The Presentation layer also handles encryption and protocol conversion functions.

Numerous protocols reside at the Presentation layer:

- ◆ **Graphics formats**—Formats such as Joint Photographic Experts Group (JPEG), Tag Image File Format (TIFF), Graphics Interchange Format (GIF), and Bitmap (BMP) handle the presentation and display of graphic images.
- ◆ **Sound and movie formats**—Formats such as QuickTime, Moving Picture Experts Group (MPEG), Windows Media File (WMF), Digital Video Express (DIVX), and RealAudio (movie) and Windows Audio Volume (WAV), Musical Instrument Digital Interface (MIDI), and Moving Pictures Experts Group Layer-3 Audio (MP3) (sound) provide for translating and presenting sound and video files.
- ◆ **Network redirectors**—Some of the most overlooked protocols that function at the Presentation layer are the network redirectors, handling the protocol conversions from your network-based formats—that is, Server Message Block (SMB) and Netware Core Protocol (NCP)—and the end user applications themselves.

Session Layer

The Session layer is responsible for setting up the logical communications channels between network hosts and applications. Each time two systems communicate, they establish a “session” that allows the hosts to differentiate between hosts and applications. The reason for this is simple—most hosts run multiple applications and are communicating between multiple hosts at the same time.

By providing a mechanism for setting up, maintaining, and tearing down the session, a single host can have multiple sessions in use while ensuring that each application (or multiple conversations occurring with a single application) keeps its data separate from any other applications. For example, if I am going to two different Web sites, I want the content for site one to appear in the browser for site one and the content for site two to appear in the browser for site two. The Session layer ensures that, even though I may be using a single application (in this case a Web browser), the data from multiple sources stays separate. Some examples of Session layer protocols include the following:

- ◆ **Network File System (NFS)**—Used with TCP/IP and Unix for remote access to resources
- ◆ **Remote Procedure Call (RPC)**—A client/server redirection mechanism (commonly used in Microsoft network environments) allowing for procedures to be created on clients (for example, the Microsoft Workstation service making a get file request) and executed on servers (for example, the Microsoft Server service handling the request and retrieving the file).
- ◆ **Structured Query Language (SQL)**—SQL provides the mechanisms for a user to access and define his or her information requirements, typically when connecting to a database.

Transport Layer

The Transport layer is primarily responsible for handling the end-to-end communications between host systems. One of the ways this occurs is via a process known as *segmentation and reassembly*. The Transport layer takes the data received from the upper layer protocols and breaks it into *segments* that are sized in accordance with the maximum segment size of the network in question. Because the data segments may arrive at the destination out of order, these segments are labeled so that the receiving system knows how to put them back together to re-create the appropriate upper-layer data. This logical communications between hosts is sometimes referred to as *virtual circuits*. Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) are two protocols that reside at the Transport layer. They will be discussed in more detail during the discussion of TCP/IP.

Network Layer

The Network layer is primarily responsible for the logical addressing of packets and the routing of data between networks. All hosts in a network are either *local* or *remote*. Local hosts are defined as those that can receive the physical signal that the source host transmits. In order to do this, hosts must share the same piece of wire. However, not all hosts do this. Sometimes the source host and destination host are in physically different locations or on physically different networks. These hosts are known as remote hosts because they cannot receive the physical signal that the source host transmits. To address this issue, and still allow the hosts to communicate, the Network layer uses logical addresses to logically define hosts so that they can be located regardless of physical location. This process of transmitting data regardless of physical location is known as *routing*. The Network layer also handles the translating of physical addresses to logical addresses. Segments that are received from the Transport layer are encapsulated within a Network layer header to become packets. Some of the protocols that reside at this layer include the following:

- ◆ **Internet Protocol (IP)**—There are some who would contend that IP *is* the Network layer. IP handles the logical addressing of hosts and the routing of data via a hierarchal addressing scheme. The benefits of a hierarchal addressing scheme are one of scaling, in that it can handle many more addresses than a flat system. In addition, it is much easier to enable routing because multiple networks can be grouped together and treated as single entries in the routing table making routing much more efficient. In fact, it would probably be impossible to route on a global scale with a flat addressing system. IP is defined in RFC 791.
- ◆ **Internet Packet Exchange (IPX)**—Used primarily on Novell-based networks, IPX provides for the logical addressing of hosts via network and host addresses.

Data Link Layer

The Data Link layer is primarily responsible for the physical addressing of frames and the translation of packets from the Network layer into bits for the Physical layer to transmit.

NOTE

The Role of Routers and Layer-3 Switches at the Network Layer

Routers and layer-3 switches are Network layer devices. They are considered Network layer devices because of the special capabilities, namely the routing of packets, that they perform. Because routers and layer-3 switches know the difference between networks, they can be used to separate broadcast domains. This simply means that routers will not forward broadcasts from one network to another network by default.

NOTE

The Effects of Broadcasts and

Collisions Broadcasts and collisions can greatly degrade network performance. Broadcasts are defined as data that is addressed for all hosts, regardless as to whether the destination can actually do anything with the data. Every host must process the broadcast, at least until it determines that the data is not for it. As a result, broadcasts can degrade performance by hampering a device's capability to transmit data because it is busy processing broadcast traffic. A common misconception is that broadcasts are more traffic than unicasts (packets that are destined for a specific host). This is simply not true. Similar to the question of "What weighs more, 1,000 pounds of lead or 1,000 pounds of feathers," a 1,000-byte broadcast is no more or less traffic than a 1,000 byte unicast. The problem lies in the amount of devices that must process the data, thus preventing them from performing other tasks. All the devices that receive the same broadcasts are known as being in the same broadcast domain. To optimize network performance, you can use routers to separate broadcast domains. In doing so, you will reduce the number of systems that have to deal with any given broadcast, thus increasing overall performance.

Collisions occur as a result of multiple devices sharing a single segment of cable. The cable can carry only a single signal at a time. The more devices that are on a segment, the greater the likelihood that two devices attempt to communicate at the same time—thus causing a collision. Collisions degrade performance by causing devices to retransmit data until they are successful.

continues

When the Data Link layer protocols receive packets from the Network layer, they are encapsulated with datalink header and footer information to become *frames*. The Data Link layer also ensures the error-free delivery of data by using a CRC (Cyclic Redundancy Check) in the frame footer. This is simply a calculation of the size of the frame prior to transmitting. The Data Link layer uses the hardware address to identify the source and destination devices. When the destination host receives the frame, it performs the check again to make sure that the value that the source host came up with in the frame footer is the same value that the destination just calculated. If it is not, the destination knows that the data sent was in error and discards it. The following protocols are among those used at the Data Link layer:

- ◆ **Institute of Electrical and Electronics Engineers (IEEE) 802.2**—Sometimes called the *LLC sublayer*, this protocol defines the interface between the Network layer and the underlying network architecture. It also provides the flow control and sequencing of control bits.
- ◆ **IEEE 802.3**—Sometimes called the *Media Access Control (MAC) sublayer*, this protocol defines how the packets are transmitted on the media. This is the point at which encapsulation of packets to frames occurs. It also provides the error checking and ordered delivery of frames.

Switches and bridges are datalink-layer devices. They are considered Data Link layer devices because of the special capabilities, namely the ability to identify the physical location of hosts, that they perform. As a result, switches and bridges can be used to segment a network while still enabling hosts to physically communicate. This reduces collisions by separating collision domains.

Physical Layer

The Physical layer is primarily responsible for sending and receiving data. The data is transmitted as bits—1s and 0s. The Physical layer also handles the specifications for the electrical, mechanical, and procedural components of the communications media. The Physical layer also identifies the DTE (Data Terminal Equipment) and the DCE (Data Circuit-Terminating Equipment) used in physical signaling and transmitting and receiving of data.

Hubs and repeaters are considered physical-layer devices. This is because they simply receive, reamplify, and forward the signal without actually looking at the data that is being transmitted.

OSI Summary

The OSI model provides a logical blueprint that can be used to understand how networking communications takes place. The OSI model is separated into seven layers. Each layer is responsible for specific tasks and functions, and for interfacing with the layer above and below itself. This modular design provides for the ability to change functions within any given layer without impacting the function of any other layers.

As data is passed down the layers of the OSI model, the data is encapsulated by the lower layer becoming segments at layer 4, packets at layer 3, frames at layer 2, and finally bits at layer 1 which are ready to be transmitted. When the destination receives the bits it simply reverses the process, unencapsulating the frames, then the packets, then the segments, eventually presenting the original data to the application that needs it.

Figure 2.2 illustrates the encapsulation process as it relates to the OSI model. Upper-layer data is received from the host and processed by the top three layers of the OSI model. At the Transport layer the upper-layer data is encapsulated with a Transport layer header and becomes known as a *segment*. The segment header contains information such as the application ports that are in use. The segment is passed down to the Network layer, where it is encapsulated with Network-layer header information and becomes known as a packet. The packet header contains information such as the transport protocol that was in use, as well as the logical source and destination addresses. The packet is passed down to the Data Link layer, where it is encapsulated with data link header and footer information to frame the packet. At this point, the data is known as a *frame*. The data link frame header contains information such as the Network layer protocol that is in use, as well as the physical source and destination addresses. Finally, the frame is turned into bits, which are then transmitted across the wire.

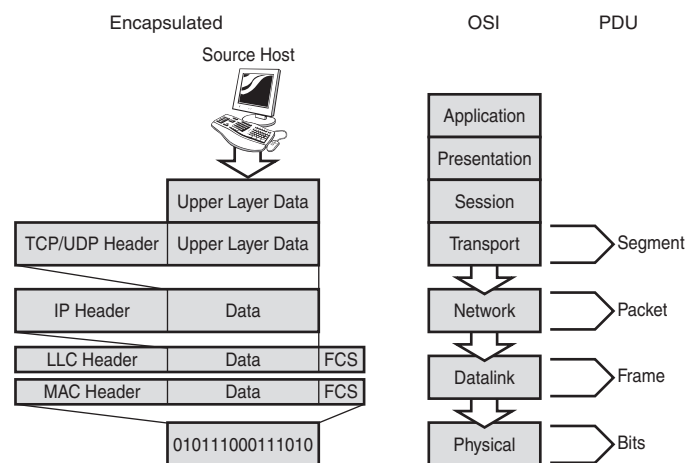
REVIEW BREAK

NOTE

continued

The devices that are capable of having their signals collide with each other are known as being members of the same collision domain. To optimize performance, you can use switches to create collision domains. In doing so, you will reduce (and potentially eliminate) the likelihood of a collision occurring, thus ensuring that the hosts need to transmit data only once, increasing overall performance.

FIGURE 2.2
Encapsulation and the OSI model.



After the destination receives the bits, it is able to rebuild the frame and process the data link header and footer to determine to which Network layer protocol to pass the data up. At the Network and Transport layers, this process is repeated using the appropriate header information, until the data can be delivered to the appropriate application. The encapsulation process allows the destination host to know what to do next with the data it is receiving and processing.

By understanding the OSI model, and understanding the layer at which specific protocols function, a security professional can understand the impact and function these protocols will have on their security design. The OSI model is like a dictionary that provides the words and language that the networks of today speak.

NETWORK CHARACTERISTICS AND TOPOLOGIES

Now that we have discussed the layered approach to networking, we can apply this information to understanding how communications and more importantly the securing of communications occur. As one would expect, different protocols and networks will require different methods of securing them. By understanding how networks function, we can determine the most effective way to secure those technologies.

IN THE FIELD**USING COMMON ATTRIBUTES**

Virtually all networks share common attributes. This is a bit of a double-edged sword. On one hand, it means that security professionals have a bit of a guideline in terms of how technologies function and how to deploy those technologies. On the other hand, it also means that malicious users have a bit of a guideline in terms of how to exploit those guidelines.

An example of this is something as simple as using a naming convention. Most good administrators use some kind of naming convention to name servers and resources so that it is easy to manage the devices and figure out what any given device does. For example, if I see a server with the name “DC01,” I can usually make a safe bet that the server is a Microsoft Domain Controller. Unfortunately, malicious users also know this, which can make it easy for them to figure out what systems to target. Does this mean that you should abandon your naming convention? Probably not, but it does mean that you need to be aware of the risks from a security perspective. This same line of thinking should be applied to understanding how network technologies function.

In keeping with “thinking layered,” we are going to begin this objective of the domain by starting with the physical-layer technologies and characteristics.

Ethernet Local Area Networks (LANs) typically utilize three types of cabling—coax, unshielded twisted pair (UTP), and fiber optic as well as wireless transmissions. Thin coax, or 10BASE-2 networks use RG58/U for cabling. 10BASE-T networks utilize Category 3, 4, 5, 5E, or better cabling. Fiber networks typically use 62.5/125 micron multimode fiber (short haul) or 9 micron singlemode (long haul). Wireless tends to use radio or microwave transmission methods.

Coax

Thin coax networks, also called *thin-net* or *10BASE-2*, use coaxial cabling with T-connectors to connect to the Network Interface Cards (NICs). Thick coax networks, also called *thick-net* or *10BASE-5*, use coaxial cabling with vampire taps and AUI transceivers to connect to the NICs.

The following cable specifications exist for coax cable:

- ◆ **RG-58 /U**—Solid copper core (0.66mm or 0.695mm), 53.5 ohms.
- ◆ **RG-58 A/U**—Stranded copper core (0.66mm or 0.78mm), 50 ohms.
- ◆ **RG-58 C/U**—Military version of RG58 A/U (0.66mm), 50 ohms.
- ◆ **RG-59**—Broadband transmissions—for example, cable TV.
- ◆ **RG-6**—Higher frequency broadband transmissions. A larger diameter than RG-59.
- ◆ **RG-62**—ArcNet
- ◆ **RG-8**—Thicknet, 50 ohms

Coax is a bus network, where all nodes communicate on a single data path, or bus. The signal on a bus network travels the full length of the bus and must stop after reaching the end of the wire. A resistor is placed at the ends of a bus system to stop the signal from bouncing back down the wire a second time. These resistors are called “terminators” and are required at each end of a bus network.

The cabling for coax has a resistance of 50 Ohms at three feet or more which means to stop a signal without bounce at the end requires a resistance of 50 Ohms. Coax networks are less commonly used than 10BASE-T networks because coax has a single point of failure for the entire segment (the line between two terminators) and is more difficult to troubleshoot.

10BASE-2 Specifications

The maximum number of nodes per segment (between repeaters) on a 10BASE-2 segment is 30. The maximum length of a segment is 185 meters. You can actually determine the maximum cable length by the name 10BASE-2. 10 stands for 10 Mbps. *BASE* stands for *baseband*. 2 stands for 200 meters (okay, so it is a little short). 10BASE-2 adheres to the 5-4-3 rule. This simply means that you can have a maximum of five segments connected via four repeaters but only three segments can have hosts on them. The two segments that cannot support hosts are called Inter-repeater Links (IRL).

10BASE-2 uses BNC (British Naval Connector) type connections for interconnectivity. There are a number of other components used for 10BASE-2 connections. The BNC connector is typically placed at the ends of each segment and connects to either a barrel connector (used for joining segments) or a Tee connector (used to connect to PCs and hubs). Figure 2.3 illustrates some of the more common 10BASE-2 connectors.




FIGURE 2.3
10BASE-2 connectors.

10BASE-5 Specifications

10BASE-5 uses a Vampire tap and a transceiver to connect to PCs and other network devices. The Vampire tap works by surrounding the cable, opening up the outer jacket shielding, drilling a hole to the conductor, and using a center probe to provide conductivity. 10BASE-5 supports a maximum of 100 taps. The transceiver provides for the connectivity to devices via Attachment Unit Interface (AUI) connections (DB15). 10BASE-5 supports a maximum of 1024 hosts per segment. The maximum segment length for 10BASE-5 is 500m. 10BASE-5 adheres to the “5-4-3” rule. This simply means that you can have a maximum of five segments connected via four repeaters but only three segments can have hosts on them. The two segments that cannot support hosts are called inter-repeater links (IRL). 10BASE-5 uses barrels and terminators, similar to 10BASE-2, but instead of BNC connectors it uses N-Type connectors. Figure 2.4 illustrates some of the common 10BASE-5 connectors.




FIGURE 2.4
10BASE-5 connectors.

Troubleshooting coax networks generally involves finding a break in the segment. The problem with coax breaks is that if a segment has a break in it, it and all the segments it is connected to are down. This includes all the computers as well. The breaks in coaxial LANs most often occur at the connectors (T connectors or barrel connectors) or at the terminators.

A Time Domain Reflectometer (TDR) can be used on one end of the cable to give an approximate distance within a few feet or so to the break in the wire. This device sends a signal similar to sonar down the cable and then times its return. The signal will “bounce” when a break is detected. The time it takes for the signal to return will indicate the distance to the break. A TDR can be found on high-quality protocol analyzers. If the failure is detected at the terminator, a Digital V-Ohm meter can be used to measure the resistance from the center pole of the terminator to the casing. The resistance should be very close to 50 Ohms.

Unshielded Twisted Pair

The most common type of cabling for Ethernet LANs is unshielded twisted pair (UTP). UTP cable comes in 10BASE-T and 100BASE-TX media type. The 10 and 100 refer to the speed the network runs at, either 10 or 100 Mbps. The cabling specification for this topology is known as Category 3, 4, 5, 5E, 6 and 7. The category of cabling indicates the quality of the signal carrying as well as the number of wires used and number of twists in the wires. These factors contribute to greater potential speeds depending on the category of cable. Table 2.1 details the different category and speed ratings of cable.

TABLE 2.1

UTP CABLE CATEGORIES AND SPEEDS

<i>Category</i>	<i>Speed Rating</i>
Category 3	Rated for voice and data up to 10Mbps/16MHz
Category 4	Rated for voice and data up to 16Mbps/20MHz
Category 5	Rated for voice and data up to 100Mbps/100MHz
Category 5e	Rated for voice and data up to 1000Mbps/100MHz
Category 6	Rated for voice and data up to 1000Mbps/250MHz
Category 7 (proposed draft)	Rated for voice and data up to 10000Mbps/ 600MHz is expected

Category 5 takes up the bulk of our discussion because it is the most widely used cable type in the industry, although most new installations should be using Category 6 or 7 cabling at this time.

IN THE FIELD

PLANNING FUTURE CABLING NEEDS

Many vendors are offering cables that are better than IEEE's CAT5. These are sometimes advertised as "Category 5 Enhanced or CAT5E, Proposed Category 6." CAT5 is all you'll need for 100Mbps Ethernet speeds, but if you want Gigabit Ethernet capabilities over copper in the future, consider cabling such as CAT5E or CAT6 that can handle 350MHZ or better.

Category 5 cabling uses RJ-45 connectors to plug into a hub, modular jack, punch-down block, or switch. Figure 2.5 illustrates some of the more common RJ-45 connectors.




FIGURE 2.5
UTP connectors.

UTP Specifications

10BASE-T is also commonly referred to as *unshielded twisted pair* (UTP) cabling. This is simply because the cable has no shielding, and the four pairs of conductors twist around each other inside the cable jacket. Because there is no shielding, UTP is very susceptible to electromagnetic interference (EMI) such as the EMI given off by fluorescent lights. As a result, UTP should not be used near such EMI sources. UTP is also very easy for a malicious user to capture the data being transmitted without ever needing to tap into the cable. Instead, such a user can run a tool that will capture the electric signal being produced, and read the data that way.

UTP has a maximum cable length of 100 meters and a maximum of four repeaters between end stations. Hubs act as repeaters. There can be a maximum of 1024 stations per network.

Troubleshooting UTP

Troubleshooting UTP is much easier than troubleshooting coax. This is one of many reasons why UTP has displaced coax as the network cabling of choice. Because UTP only supports two devices on a cable (that is, a computer and a hub or switch), when a cable failure occurs it is generally easy to pinpoint. As you will see in cable plant design in the next section, much of UTP troubleshooting is simply tracing the cable back to the source. Using a TDR with coax can assist this, but generally it takes longer to set the TDR up than it does to just follow the cable back. Some common culprits with UTP problems are using incorrect patch cables and incorrectly crimping/punching down the cable (discussed more in the next section). Generally however, if you have a link light with UTP the problem is somewhere else.

Fiber Optic

Fiber-optic cable is predominately used for backbone and device interconnectivity as opposed to end user connectivity. There are a couple of reasons for this. First, fiber is much more expensive than UTP or coax. Additionally, because fiber-optic cable is made of glass, it is much more fragile than the alternatives. Let's face it—we all know what our users' work environments look like. Fiber doesn't stand a chance! That's okay though, because fiber has a role to which it is much better suited—device interconnectivity on the backbone. Fiber has now replaced 10BASE-5 as the predominant backbone device interconnectivity method. This is due to the speed and distance at which fiber optics can transmit.

Fiber-Optic Cable Components

Fiber-optic cable is made of a buffer, usually PVC or rubber, and the actual fiber. The actual fiber strand consists of two pieces of fiber. One is called the *core* and the other is the *cladding*. The core is the propagation path for light and the cladding, which has a different density than the core, acts as the refractive layer. The core is made of silica glass or plastic ranging in size from 8 microns (μm) to 1000 microns. The cladding reflects the light that tries to escape the core so the light stays in the core.

A coating (also called a *buffer*) surrounds the cladding. In a tight buffer construction, the buffer is directly on the fiber. In a loose buffer construction, there is a layer of gel between the buffer and the fiber. This constitutes a single strand or piece of fiber. Figure 2.6 illustrates the components of a piece of fiber cable.

The individual fiber strands are then typically bundled in pairs, or multiple pairs, because each fiber can only send a signal in a single direction. A reinforcing layer of plastic (the outer jacket) is placed around the individual strands. The strands are also wrapped in Kevlar to provide both strength as well as flexibility to the actual fiber strands further reinforcing the fiber-optic cable.

One-pair fiber cable, which is typically used in patch cord implementations, is generally called *simplex* or *zipcord*. Multi-pair fiber cable that is double buffered (tight buffer with outer jacket) is generally referred to as *distribution cable*. Distribution cable does not reinforce the fibers, and thus to terminate the cable one needs to use a *breakout box*. A breakout cable is made of several simplex/zipcord cable bundles and is generally more rugged because the fiber can be terminated like zipcord (because it effectively is just a bundle of zipcord). Loose tube cables are composed of several fibers together in a plastic tube. The tubes are then wound around a central strength member and jacketed providing a high fiber count (in the 100s). The tubes are filled with gel to prevent harm and protect the buffer, which is very thin. While this cable must be handled carefully, it is well suited for outdoor and very large backbone (that is, service provider) implementations.

Breakout kits are used for terminating fiber in a loose buffer tubes. In a loose buffer tube construction, the fiber is contained in a gel-filled polymer tube that has an inner diameter larger than the fiber itself. This provides a high level of isolation for the fiber from external mechanical forces. A loose buffer is used in outdoor applications and can accommodate the changes in external conditions (that is, contraction in cold weather and elongation in warm weather).




FIGURE 2.6
Fiber cable components.

Multi-Mode Fiber

There are two main types of fiber-optic cable, *multi-mode* and *single-mode*. Multi-mode fiber is mainly used for short or medium distances and for low bandwidth applications. The actual fiber sizes used (core/cladding) are

- ◆ 50/125 μm
- ◆ 62.5/125 μm (most common)
- ◆ 100/140 μm

It is called multi-mode fiber because the fiber is designed to carry multiple light rays, or modes, concurrently, each using a slightly different reflection angle within the fiber core. The modes disperse over longer lengths (this is called modal dispersion) and is one of the reasons multi-mode is suited to shorter distances. For 100Mbps Ethernet, the distance limitation is 2km. For 1Gbps Ethernet, the distance limitation is around 550m.

Single-Mode Fiber

Single-mode fiber is designed for the transmission of a single ray, or mode, of light as a carrier and is used for long-distance communications. Because there is only one ray of light, a smaller core can be used for single mode fiber. The actual fiber sizes used (core/cladding) are

- ◆ 8/125 μm
- ◆ 9/125 μm
- ◆ 10/125 μm

Single-mode fiber can achieve much greater distances than multi-mode. For 100Mbps Ethernet the distance limitation is 20km or more! For 1Gbps Ethernet, the distance limitation is about 3km for long haul while extended distance transmission can reach up to 100km.

Fiber connectors come in many types. The most commonly used connectors are the Stick and Turn (ST), Stick and Click (SC), and SC Duplex connectors. Fiber is attached to the connectors via splicing. There are two main types of splices, fusion and mechanical.

Fusion splices use a welding process to fuse the fiber to the connector (or to other pieces of fiber). This provides a stronger and lower-loss connection. Mechanical splices use an alignment fixture to mate the fibers and then you either polish the end of the fiber (very hard to do and time consuming) or use a matching gel or epoxy (more common) to minimize the reflection.

Dense Wave Division Multiplexing

Dense wave division multiplexing (DWDM) is one of the newest forms of fiber-optic transmission. DWDM works by the principle that different color light resides at different frequencies and the light at one frequency does not interfere with light in a different frequency. Think of how a prism works. A prism can be used to break the individual colors of light out. DWDM does the same kind of thing, taking different colors of light and breaking them out at the source and destination. The advantage is that you can have multiple *channels* of data being transmitted simultaneously without impacting the throughput of any channel. Currently from 4 to 32 channels of wavelength are supported, but that number is increasing even as I write this, with future expectations of 80–128 channels. For example, an OC-48 fiber transmits at 2.5Gbps. Using four channels, the speed can be increased to $4 \times 2.5\text{Gbps}$, or 10Gbps with no new fiber needed. At 32 channels, the throughput is 80Gbps. Future implementations then would deliver 320Gbps of data on a 2.5Gbps OC-48 link. But hey, why stop there? Consider OC-256 at 13.271 Gbps. Using 128 channels, we would have a bandwidth of 1.699 Tbps. That's terabits per second. Quake-fest, here I come! Obviously, this technology plays well in environments where all the fiber is in use but more connections are needed, as well as in oceanic cables because smaller cables can be used to deliver the same data it used to take a cable 10 times as large to do.

Wireless

Wireless is finding its way into more and more networks for a very simple reason—because there are no wires, the devices can be located anywhere that they can receive a signal. A big push for wireless has been with the small office/home office (SOHO) users, because many houses were not designed with network cabling in mind.

88 Part I EXAM PREPARATION

By simply using a wireless network, the user can place the computer (or multiple computers) anywhere in their home and still have network access. In corporate environments wireless is often used in executive and campus environments. This allows executives to travel anywhere on the executive floor and still be able to access the network without needing to reconfigure or recable anything. Another increasingly popular deployment of wireless has been with Point of Sale (PoS) systems. Rather than running cabling to all of the systems handling the sales transactions, they simply run wireless.

There are a few rather substantial drawbacks to wireless at this time. The first is the lack of standardization, or more appropriately the fact that there are numerous incompatible and competing standards being employed. From 802.11 Wi-Fi to 802.11a to 802.11b to 802.11g to 802.15 Bluetooth, wireless standards definitely live by the “the greatest thing about standards is there are so many to choose from” adage. The thing to remember is to make sure that all of the equipment you select supports the same standard.

The other problem with wireless is one of security. In the same way that anyone can tune a radio to receive certain radio stations, people can connect to a wireless network by simply running the appropriate equipment and being within a certain range. This makes it easy for malicious users to compromise a system, and in fact fairly recently a certain chain of stores found itself in a bit of a problem when it was discovered that its PoS systems ran wireless with no security, so anyone sitting in the parking lot with a wireless card and a laptop could potentially be capturing credit card transactions. Another drawback is that interference can severely limit distances that wireless networks cover.

The lesson to be learned here is to secure your wireless environment using authentication and encryption. The authentication ensures that only authenticated devices can connect to the network and the encryption will ensure that even if intruders can capture the signal, they must decrypt it to gain any data of value.

NETWORK TOPOLOGIES

Virtually all networks use one of the following topologies:

- ◆ Linear bus
- ◆ Star
- ◆ Ring
- ◆ Tree
- ◆ Mesh

We explore the different network topologies in more detail in the following sections.

Linear Bus Topology

One of the earliest networking topologies was the linear bus topology. In a linear bus, all the systems were connected in a row to a single cable in a *daisy-chain* fashion. This simply means that the cable runs from system 1 to system 2 to system 3. The piece of cable that all the systems were connected to is known as a *segment*. Coax-based networks were classical physical linear bus topologies, while Ethernet is a classical logical bus topology. We will talk more about Ethernet in a little bit. Figure 2.7 illustrates how a linear bus topology is connected with all the computers sharing a single piece of wire.




FIGURE 2.7
Linear bus topology.

Understanding how devices communicate on a linear bus requires an understanding of three core concepts used in linear bus networks:

- ◆ How the signal is transmitted
- ◆ Signal bounce
- ◆ Signal termination

Physically, the signal is sent to all devices connected to the linear bus segment. On the surface, this may sound like the signal is a broadcast, but that is not the case. Instead, this is simply a matter of electronics and electricity. If I take a lamp cord and cut the jacket off the cable, this exposes the conductor. If I then have a bunch of people grab a hold of the cable and plug it into the wall, they are all going to get shocked. A linear bus works in the same fashion. When the devices are connected to the bus, they all share a common conductor, which means that when an electric signal is put on the wire (for example, during data transmission) all the devices connected to the segment are going to get the electric signal. This does not mean that all the systems actually process the data. We talk about this more when we look at Ethernet and switches.

Another thing to understand is that only one signal can exist on the segment at a time, which means that only one device can transmit at a time. As a result, the more devices that you connect to a linear bus, the worse the performance degradation will become. This is known as *contention*, which simply means that the devices are in contention for the same segment to transmit. A linear bus is also known as a *passive technology* because the devices on the segment do not move the data from one device to the next; rather the signal is generated at the source and all other devices passively receive the signal.

As the signal is put on the wire and begins to move away from the source, it encounters the problem of signal bounce. After the signal hits the end of the cable, the signal bounces back and continues to travel back and forth, effectively preventing any other systems from being able to communicate. In order to address this, a linear bus uses terminators at the ends of the bus to absorb, and thus terminate, the signal. The logic behind this is really quite simple.

By the time the signal has reached the terminator, every other device on the bus should have been able to receive the signal and either process or discard the data accordingly.

One of the problems with a linear bus has to do with termination. If any part of the bus is not properly terminated, the entire bus will cease to function properly. From a security perspective, this means that someone can take out all of the devices on the bus by simply removing the termination (for example, by cutting the cable). Linear bus is very susceptible to cable faults as a single point of failure.

Star Topology

Unlike coax, the topology method in a 10BASE-T network is a *star* because all devices must have a segment of wire connecting them to an active hub or switch before being capable of communicating with other devices on the LAN. In other words, each computer effectively has its own piece of cable with the computer on one end and the network device on the other. Figure 2.8 illustrates a star topology with all the computers connected to a central hub/switch.




FIGURE 2.8
Star topology.

The benefit of this type of system is when there is a cable fault only the device on that cable is affected, unlike coax where all devices on the segment are affected. Logically, however, a 10BASE-T network still operates as a bus. So although each computer is on a different physical cable, all the computers are logically connected as a linear bus due to the hub/switch.


Star topologies are also used to implement what is known as a *collapsed backbone*. In a traditional network, the backbone of the network consisted of cabling running between multiple network connectivity devices (often in a linear bus fashion). The collapsed backbone replaces this by having the network devices connected to a single device that actually provides the backbone connectivity. Because a collapsed backbone requires less cabling, it is considered cheaper and easier to maintain than traditional backbones.

The network is not affected by individual cable faults because the hub/switch will short the port on which a cable fault occurs, effectively closing the linear bus and allowing the other devices on the network to continue functioning. However, because the hub/switch is the center of the star, it becomes a single point of failure, because if it stops functioning the devices can no longer communicate with each other.

The star topology has become the most used network topology today.

Ring Topology


The ring topology is designed using a loop of cable to interconnect the devices. The signal is transmitted in a single direction around the loop, with each device retransmitting the signal as they receive it. The ring topology is considered an active topology, unlike the linear bus, because of this. One of the drawbacks of this type of system is that if any system stops passing the signal, or starts generating bad signals, it can take the entire ring out. Figure 2.9 illustrates the design of a ring topology.

**FIGURE 2.9**

Ring topology.

Tree Topology

The tree topology is based in part on the bus and the star topology. In the tree topology devices are interconnected to each other via bus connections; however, there are multiple nodes supported on each potential branch, as shown in Figure 2.10.

**FIGURE 2.10**

Tree topology.

Mesh Topology

The mesh topology (sometimes called the mess topology) ensures that every node on a network is connected to every other node.




FIGURE 2.11
Mesh topology.

Mesh networks are typically deployed to create backbone and WAN networks. In a full mesh topology, all nodes are connected to each other. In a partial mesh, multiple full mesh networks are interconnected to each other, though every node does not necessarily connect to every other node. Figure 2.11 illustrates a full mesh topology.

LAN and WAN Technologies

As mentioned, virtually all networks use one of the previously mentioned physical topologies. The various LAN and WAN technologies build upon the topology to provide an effective method of sending and receiving data. Although the topology may stipulate that the signal is generated and all hosts receive it, it is the role of LAN and WAN technologies to figure out what a device actually does when the signal is received.

Data is transmitted on LANs using one of three transmission techniques:

- ◆ **Unicast**—The packet is addressed to a specific destination host, both physically and logically.
- ◆ **Broadcast**—The packet is destined to all hosts on a subnet or network. At the Data Link layer the address used is FFFFFFFF in hexadecimal. At the Network layer the address used is the network broadcast identifier or the all-networks broadcast address of 255.255.255.255. There is a variant on broadcasts known as a *directed broadcast*. In a directed broadcast, the Data Link layer destination address is a broadcast, but the Network layer destination address is a unicast address. ARP is sometimes referred to as a directed broadcast.
- ◆ **Multicast**—The packet is addressed to multiple hosts via the use of group membership addresses. Multicasts play the middle ground between needing to repeatedly send unicasts to multiple destinations and broadcasting to all destinations, even though only a subset of the hosts needs the data. With a multicast, the data is sent only to the systems that register as wanting it, thus reducing the overhead of a broadcast and the excess packets that would be needed to transmit via repeated unicasts.

Ethernet

Ethernet is the single most predominant technology in use today. With speeds ranging from 10Mbps to 10Gbps, Ethernet possesses awesome speed and scaling capabilities. Today, most Ethernet is physically cabled as a star topology, but remember that logically it still functions as if it were a linear bus. This means that all Ethernet devices expect communication to occur as if they were connected to the same physical cable segment.

Ethernet is specified in the IEEE 802.3 specification as a Carrier Sense, Multiple Access/Collision Detection (CSMA/CD) methodology.

Ethernet networking is known as a contention-based media access methodology that allows all hosts on a network to share the same bandwidth of a link. The problem is that only one host can be transmitting or receiving on a link/segment at any given time. Think of contention in terms of a busy house with a single bathroom in the morning. Everyone needs to get in and shave, shower, brush their teeth, and so on, but there is only one bathroom and only one person can be using it at a time. As a result, everyone is in contention for the use of the bathroom. The fewer people who need to use the bathroom at any given time—for example, during the summer when the kids are sleeping late—the faster it is for everyone to get in and out. Once school starts up again though, everyone winds up spending more time to do the same tasks, because they have more people to wait on and share time with before they can make use of the bathroom.

To address the contention inherent to all Ethernet implementations, Ethernet uses CSMA/CD. This helps the devices on the network share the bandwidth while making sure that two devices cannot use the bandwidth at the same time. The problem when two devices attempt to use the bandwidth at the same time is the creation of collisions. When two hosts attempt to transmit at the same time, they both generate a signal and place the signal on the wire. A basic rule of conductivity is that only one signal can be carried at a time, and thus when the two signals meet, the data they are carrying “collides” causing the data to be lost. The use of CSMA/CD is also known as collision management, because it helps to eliminate collisions from occurring. So how does it do that?

NOTE

802 Standards on the Web The IEEE recently made the entire 802 standards documentation available for free online. You can now download the standards from <http://standards.ieee.org/getieee802/portfolio.html?agree=ACCEPT>. Although reading standards is not exactly the most pleasurable reading experience, there is no substitute if you truly want to know how things work.

CSMA/CD works like this: When a host wants to communicate on the network, the host listens to the wire to see whether it detects a signal. If it doesn't detect a signal, the host attempts to transmit. If it does detect a signal, the host waits and then checks again. This is similar in concept to pulling out onto a road. When you want to get on the road, you look both ways to see whether there are any cars. If you see cars, you wait. If you don't see any cars, you begin to pull out onto the road. This is the *Carrier Sense* portion of CSMA/CD.

Ethernet hosts also realize that even though they detected an open signal, all of the other hosts are sharing the bandwidth with them. As a result, multiple devices could be accessing the networks (the *Multiple Access* portion of CSMA/CD). After sending data, the host then listens to the wire to see whether any other hosts attempt to transmit at the same time, or during their transmission cycle. If they detect a signal from another device, they send out a signal to jam the media that causes all of the hosts to stop transmitting and wait a random time period before they begin sending their data again. This is similar in concept to you continuing to look around even after you start pulling out on the road in case you missed something, or in case someone suddenly shows up. If they do, you honk your horn to tell them that there is a problem and we need to wait and let one car go before the other.

Even with all these precautions though, cars still crash and packets still collide. To address this, Ethernet uses Collision Detection (the last part of CSMA/CD) to detect whether there was a collision so that the hosts that were transmitting know that they need to retransmit their data, because a collision occurred. In a sense, this is kind of like the insurance company after a crash fixing your car so that you can start driving it again.

Ethernet also has the ability to function in either *half-duplex* or *full-duplex* mode. Half-duplex is what was originally specified in the 802.3 Ethernet specification and uses a single wire pair with data running in both directions on the wire. As a result, hosts can either transmit or receive, but they can't do both at the same time. It is kind of like using a walkie-talkie. One person talks and the other person listens. The other person can't start talking until the first person is done. Ethernet uses CSMA/CD to try to minimize collisions, but that doesn't always work. As a result, it is a common statement that Ethernet in half-duplex mode is only 50%–60% efficient.

Chapter 2 TELECOMMUNICATIONS AND NETWORK SECURITY 97

This is a bit of a misstatement, or at least an oversimplification, as there are many other variables that can adjust that figure up or down such as number of hosts, frequency of traffic, and so on.

On the other hand, full-duplex Ethernet uses two pairs of wires, instead of one pair like half-duplex Ethernet does. Full-duplex is a point-to-point connection between the transmitting and receiving hosts. Full duplex also allows the devices to send and receive at the same time, because two paths can be used—a transmit path and a receive path. Because of the nature of full-duplex Ethernet, the old rules of CSMA/CD are changed a little. First, in order to run full-duplex Ethernet, each host must be plugged directly into the switch with no other hosts on the segment. In fact, full duplex is only available in switched environments. Because there isn't anyone else on the segment by definition, the Multiple Access part of CSMA/CD can be eliminated. Also, because there isn't anyone else on the segment, the host doesn't need to Carrier Sense to see if there is a signal. That leaves us with CD, which we also don't need because transmit and receive use physically separate wire pairs, and thus there is not a possibility (never say never though) that a collision will occur.

Because of the elimination of the overhead of CSMA/CD, full-duplex Ethernet is generally regarded as upwards of 100% efficient. Likewise, full-duplex Ethernet is also commonly called *200Mbps Fast Ethernet* or *20Mbps Ethernet*. This is a little bit of a misnomer. Use Fast Ethernet as an example: because it operates at 100Mbps, and we can run in full-duplex, we can send at 100Mbps and receive at 100Mbps. Due to the wonder of marketing, this is referred to as *200Mbps*, and leads some folks to believe that they are doubling the speed of their network by running in full-duplex mode. The reality is that you aren't going any faster in any one direction, sending or receiving, but you can get the same speed in both directions at the same time.

At this point you might be saying, "Okay, why doesn't everyone run full duplex everywhere?" One problem is that full duplex is not an exact science. Each vendor implements its own mechanisms for performing full-duplex operations, and often they just don't work well together. Another issue is one of scaling traffic. While it would seem like the more bandwidth we have regardless of location would be great, the reality is that it is best to scale your bandwidth from smaller pipes at the clients, to larger pipes at your servers and uplink connections.

As a result, it is generally recommended to run full-duplex on connections between network equipment, and to run full-duplex on your servers, but to allow the clients to operate in half-duplex mode only. In a way this also makes sense. Servers and network equipment need to send and receive simultaneously but clients usually don't need to do both at the same time.

Token-Ring and FDDI

The most predominant method of transmitting data on a ring topology is through the use of something called *token passing*. The token is simply a packet that data is appended to for transmission. As a result, if a system wants to transmit, it must have the token so that it can append the data to the token and transmit it. The token is passed around the ring until it arrives at the system the data is destined for, or it is received by the active monitor two times, in which case it removes the data assuming that the destination system is not online. This type of system is known as *token-ring architecture*, and operates at 4Mbps or 16Mbps speed whereas FDDI operates at 100Mbps. Token-Ring is defined by the IEEE 802.5 specification.

Token-Ring uses a logical ring, although much like Ethernet, it is primarily cabled as a physical star today. As mentioned, a ring topology tends to be an active topology, which means that the devices actively participate in the passing of data. Token ring accomplishes this by designating specific functions that ports are responsible for. *Station Ports* exist on token ring NICs and connect to the ring. *Lobe Ports* exist on the token ring hub or MAU (Multi-Access Unit) and are responsible for connecting to station ports. *Ring in/Ring out (RI/RO)* ports are responsible for connecting one ring to another ring to create a single larger ring.

Connected to the Token-Ring network, each system has a responsibility to ensure that the data is properly passed. The first responsibility is to ensure that the token is generated and exists on the ring. Without the token, no systems can send data. The responsibility of generating the token, removing bad tokens, providing clocking, maintaining ring delay, handling orphaned frames and purging the ring is bestowed to the *active monitor*. The active monitor is typically the first system brought online.

Although Token-Ring was designed to be a largely self-healing and redundant network technology, the information required to keep the ring functioning properly is something that can be used by a malicious user in regards to gaining information about how the network is designed. Additionally, if malicious users can take over the role of active monitor, they can effectively take the ability of the ring to pass data out, causing a denial of service (DoS).

Attached Resource Computer Network

Attached Resource Computer Network (ARCnet) is a largely dead network topology that was based a little bit in Ethernet and a little bit in Token-Ring. ARCnet used a Carrier Sense, Multiple Access/Collision Avoidance (CSMA/CA) access methodology to transmit data, which was based in the need to use a token in order to transmit. The catch is that ARCnet was a bus topology, not a ring. ARCnet is referred to as a token-bus network, and is the platypus of network topologies (it seemed to have bits of everything else included).

LAN DEVICES

Now that we have seen the theory and architecture with which LAN networks are built, as well as the physical interconnection methods and networking types, we need to take a look at the components and technologies that make up a network.

Today's networks are primarily made up of five categories or types of devices. Each type of device has unique capabilities, functionalities, and vulnerabilities that as a security professional, you must be aware of.

Hubs and Repeaters

Hubs and repeaters are physical-layer devices. Functionally, hubs and repeaters do the same thing; however, hubs tend to have more ports than a repeater does. As a result, hubs are sometimes called *multi-port repeaters*. I will use the term *hub* to refer to both devices.

The primary function of a hub is to receive a signal, amplify the signal, and repeat the signal out all ports. Hubs never check the integrity of the data, which means if the data contains an error, the hub will simply pass the error around. In addition, hubs do nothing to reduce contention on the network; in fact, hubs can increase contention by providing the means for more devices to connect to a segment. Because the hubs are Physical layer, there is very little that can be done to secure traffic or devices connecting to hubs. Hubs will generally pass any and all data, good, bad, and indifferent.

Switches and Bridges

Switches and bridges are datalink-layer devices, and pick up in functionality where hubs stop. Much like hubs and repeaters, switches and bridges are functionally very similar, and in fact most of these types of devices today are going to be switches. I will use the term *switch* to refer to both devices. The big differences between bridges and switches are

- ◆ Switches are hardware based and use ASICs (Application Specific Integrated Circuits) to make decisions while bridges use software. This allows a switch to function faster than a bridge.
- ◆ Switches have more ports than bridges do, and sometimes a switch is actually called a “multi-port bridge.”
- ◆ Bridges can only run one instance of spanning tree, whereas switches can have multiple instances. Spanning tree is a protocol, defined in the IEEE 802.1d standard, that is responsible for preventing loops from occurring on a bridged/switched network. Network loops at layer 2 can create a condition known as a *broadcast storm*. A broadcast storm simply means that so many broadcasts are occurring that other traffic is unable to occur. Spanning tree prevents loops by determining all the redundant paths in a network, and then blocking any paths that would create loops. This allows a network to have redundant paths; however, only one path will be available at a given time.

Switches are considered datalink, or layer-2, devices because a switch is Data Link layer aware. This means that switches understand how physical addressing occurs, and thus can use this knowledge to optimize network communications. If you recall, Ethernet networking is based on the principle that all devices share a common segment, and thus each device receives all signals all the time. This is a fundamental of networking. Unfortunately, this also means devices physically receive signals that are not destined for them which they must discard. In a sense, it is like getting a lot of phone calls for the wrong number.

Because switches understand that the signal carries data destined for a specific host, rather than just forwarding the signal blindly like a hub does, the switch will read at least part of the data and attempt to determine to which port the destination host is connected (this is known as the *destination port*). If the switch can determine the destination port, rather than forwarding the signal out all of its ports, it sends the signal only on the destination port. If the switch is unable to ascertain the destination port, then the switch falls back to *basic Ethernet* and forwards the signal to all ports. This concept is known as *segmentation*. If you recall, a segment is the cable that devices share. By intelligently sending the signal only to the ports that the destination is on, the switch effectively causes each port to be considered its own segment. Because of the number of ports that a switch can have, this is sometimes referred to as *micro-segmentation*. As a result, switches can be used to reduce the contention that is inherent to Ethernet networking which allows for a network to contain more hosts and effectively function at higher speeds.

Because switches are datalink aware, they can be used to provide some security capabilities. One of the ways that switches can do this is via the use of Virtual Local Area Networks (VLANs). The other way that they can do this is via the use of port-based security.

VLANs

VLANs are the creation of logically segmented networks within a single switch, or within a single switch fabric. A switch fabric is a group of switches that are physically connected to each other.

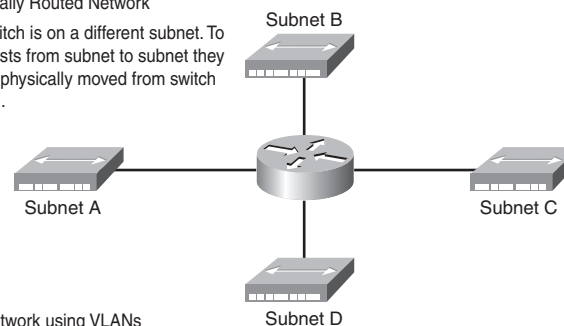
Although the primary goal of VLANs is typically the separation of broadcast domains and the creation of subnets, an additional benefit can be one of security. A VLAN is effectively a subnet, so just like a router is needed to communicate between subnets, a router is needed to communicate between VLANs. As a result, you can gain a degree of security by separating hosts between VLANs and then restricting the traffic at the router. An example of this scenario would be segmenting the HR (Human Resources) equipment on a separate VLAN from the rest of the network. This allows the administrator to control the devices that can access the HR equipment—for example, only allowing the HR workstations—by restricting traffic at the router. Figure 2.12 shows a comparison of VLAN routed and traditional routed networks.

FIGURE 2.12

Comparison of VLAN routed and traditional routed networks.


Traditionally Routed Network

Each switch is on a different subnet. To move hosts from subnet to subnet they must be physically moved from switch to switch.



Routed Network using VLANs

Each switch is a member of all 4 VLANs and thus all 4 subnets. To move hosts from one subnet to another subnet, you simply change the VLAN the port is a member of (changing the host IP address as required).



Although it might seem like VLANs are a great security mechanism, there is a drawback. Because the VLAN is logical, and the ports from both VLANs are often on the same switch or switch fabric, it is possible for data to physically transfer from one VLAN to another VLAN, even though it normally shouldn't. There have been numerous exploits, particularly with buffer overruns, that allow packets to traverse VLANs without being routed. As a result, it is generally not recommended to use VLANs when segmenting internal and externally accessible networks (for example, when using a VLAN to separate a screened subnet and the internal network).

Routers

Routers continue to build on the technologies that we have previously discussed. Routers function at the Network layer, and are often referred to as a *layer-3* device. You may have heard of layer-3 switches as well. A layer-3 switch is simply a hybrid device that combines layer-2 and layer-3 functionality, allowing the switch to forward frames when possible and route packets when needed. Because switching occurs at layer 2, it is faster than routing. As you would expect, layer-3 switches are particularly suited for VLAN environments.

Routers are able to further optimize network traffic by utilizing the logical addressing information available from the Network layer. Routers are considered “network aware” which means that routers can differentiate between different networks. Routers use this information to build routing tables, which are tables that list the following basic information:

- ◆ All the networks the router knows about
- ◆ The remote router to use to connect to those networks
- ◆ The paths, or routes, to the networks
- ◆ The cost, or metric, of sending data over the paths

With this information, the router can make intelligent determinations of the most efficient, or at least what the router deems most efficient, path to the specified network.

NOTE

Using Switches As a Security

Mechanism Port-based security is used particularly in environments that are extremely security conscious. With port-based security the administrator configures the switch to only allow a specified MAC (media access control) address to be allowed to connect. While this can provide significant security to a network, it also has the potential to require a tremendous amount of overhead. Any time that computers move or the NIC is changed, the administrator has to update the switch accordingly.

Routers are also used to segment large networks into smaller ones, as well as to reduce broadcasts on a network. Routers recognize that most broadcasts are specific to the network that they originated, so instead of forwarding the broadcast as a hub or switch does, the router will stop the broadcast.

Because routers function higher in the OSI model than switches, they are also able to provide better traffic management and security capabilities than switches or hubs can. Routers are able to examine logical addresses as well as the layer-3 header information to determine what application ports are being used and use this information for traffic filtering and blocking purposes.

Firewalls

Firewalls have achieved a status as a panacea of sorts, a generic cure all for a company's security woes. Unfortunately, firewalls—while still a great security measure—are not the be-all and end-all that some would have you believe. Instead, firewalls should be considered but a single component of a comprehensive security design.

Firewalls are designed to prevent traffic that is not authorized from entering or leaving a network. They are typically deployed as a perimeter security mechanism to screen Internet traffic that is attempting to enter the network. There are six main types of firewalls, sometimes referred to as “generations”:

- ◆ **Packet filtering**—Packet-filtering firewalls are very similar in use and function to routers. In fact, many routers include packet-filtering capabilities. Packet filtering firewalls function by comparing received traffic against a rules set that defines what traffic is permitted and what traffic is denied. This is typically performed by using IP addresses and/or port numbers to identify permitted and denied traffic. If the received packet matches the permitted traffic list, it is allowed to proceed. If it does not, the firewall discards the packet. Packet-filtering firewalls generally operate faster than other firewall types because they often do not need to read more than the layer-3 or layer-4 information in a packet before making a filtering decision. Packet-filtering firewalls are considered to be first-generation firewalls.

- ◆ **Application proxy**—Application-filtering firewalls function by reading the entire packet up to the Application layer before making a filtering decision. Whereas a packet-filtering firewall generally cannot differentiate between the valid application data and invalid application data, the application proxy firewall can. This allows an application proxy firewall to be able to recognize CodeRed data in an HTTP request, and thus block it, where a packet filtering firewall would not. Although this can provide a much higher degree of filtering capabilities, it also means that application proxy firewalls are generally slower than packet filtering firewalls. Another drawback is if a proxy does not exist for a service that the user requires, you may need an additional proxy, or you may not be able to communicate using the given service at all. An application proxy firewall is sometimes referred to as an ALG (Application Level Gateway) and is considered a second-generation firewall.
- ◆ **Circuit proxy**—Circuit proxy firewalls are a bit of a hybrid between application proxies and packet-filtering firewalls. With a circuit proxy, the firewall creates a circuit between the source and destination without actually reading and processing the application data. In that sense, it is a proxy between the source and destination. However, because it doesn't actually process the application data, it's functionally similar to a packet filter.
- ◆ **Stateful inspection**—All firewalls being considered today should perform *stateful packet inspection*. When a host sends a packet to the destination, the destination is going to process the data and potentially send a response. This network connection state is tracked by the firewall and then used in determining what traffic should be allowed to pass back through the firewall. For example, if the firewall knows that a request was sent to a Web site, because the firewall knows that the connection state is “waiting for a response” when the response comes in, rather than blocking the packet as would be normal, the firewall allows the traffic to proceed. Because these firewalls can examine the state of the conversation, they can even monitor and track protocols that are otherwise considered “connectionless,” such as UDP or certain types of remote procedure call traffic.

Stateful packet inspection can also protect against attacks that might occur as part of a normal conversation between hosts. When two hosts decide to communicate and establish a session, they define how to handle the situation of fragmented packets. Fragmented packets can occur for many reasons—for example, the original packet was too large to traverse a network segment. In those cases, the original packet could be broken down into multiple new packets by any router in the path. When the destination receives these fragments, it uses the fragmentation ID to determine the order in which the fragmented packets should be put back together to create the original packet or data. Only the first fragment contains the high layer header information that filtering decisions are made with. All subsequent fragments simply contain the necessary IP information required to properly deliver the data.

Originally, it was decided that all non-first fragments would be permitted through a filter, but the first fragment would need to match a permitted filter or it would be dropped. The logic was that without the first fragment, subsequent fragments cannot be put back together, thus the risk is minimal. Hackers realized this and found new ways to exploit networks. For example, by sending fragments that contained overlapping information, hackers found that often they could initiate a denial of service, or in some cases could even cause data to be passed between the hosts. Stateful packet inspection can deal with this by observing fragments, and only allowing fragments that it finds in the appropriate state. Additionally, many stateful packet inspection firewalls will actually perform the packet reassembly, so if the fragments contain harmful data, the firewall can reassemble and drop the data before the destination host is affected.

Stateful inspection firewalls are considered third-generation firewalls.

- ◆ **Dynamic packet filtering**—A dynamic packet filtering firewall is generally used for providing limited support of connectionless protocols like UDP. It functions by queuing all the UDP packets that have crossed the network perimeter, and based on that will allow responses to pass back through the firewall.

- ◆ **Kernel proxy**—Kernel proxy firewalls are typically highly customized and specialized firewalls that are designed to function in kernel mode of the operating system. This provides for modular, kernel-based, multi-layer session evaluation using customized TCP/IP stacks and kernel level proxies.

In truth, the type of firewall a given product is typically lies in the realm of a hybrid firewall of one or more of the six base firewall types. For example, many firewalls combine packet filtering stateful packet inspection and circuit proxy functionality all in one.

In addition to the six types of firewalls, there are four general types of firewall architectures. Some of the architectures are based on one of the firewall types, and some of them are portable design concepts:

- ◆ **Packet-filtering routers**—Packet-filtering routers are designed to sit between an internal “trusted” network and the external “non-trusted” network. Because a packet-filtering router sits along the boundary between the two networking types, it is often referred to as a boundary or perimeter router. Security is maintained through ACLs (Access Control Lists) that define the IP addresses, protocols and port numbers that are allowed. Unfortunately, maintaining the ACL can be a very complex and time-consuming process. Other drawbacks include a lack of authentication and generally weak auditing capabilities. Packet-filtering routers can provide an excellent first security boundary as a bulk filtering device due to their speed and are sometimes used to control access to a DMZ. Figure 2.13 is an example of a packet filtering firewall solution.


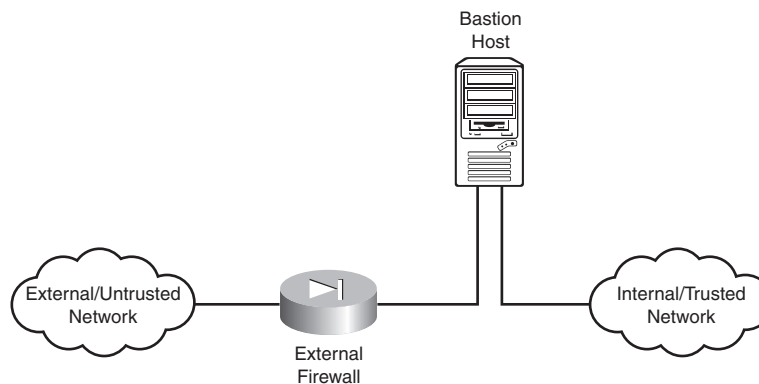


FIGURE 2.13
Packet-filtering firewall.

- ◆ **Screened-host firewall**—Screened-host firewalls typically employ both a packet-filtering firewall and a bastion host to create a firewall system. A bastion host is a system that is directly exposed to external threats. In a screened-host firewall system, the bastion host resides on the internal network, but it is the only host on the internal network that is accessible to external hosts.

This system requires an intruder bypass the external router (packet filtering) and the bastion host (proxy) in order to gain access to internal resources. Unfortunately, because the bastion host is directly connected to the internal network, if it is compromised, there is nothing to stop the intruder from having full run of the internal network. Screened-host firewalls are particularly suited to providing low-risk, limited access for connections from the Internet. Due to the lack of protection between the bastion host and the internal network, it should never be used for high-risk access such as public Web server access. Figure 2.14 is an example of a screened-host firewall system.

FIGURE 2.14
Screened-host firewall.



- ◆ **Screened-subnet firewall (with demilitarized zone [DMZ])**—Screened-subnet firewall systems provide an additional degree of network security by introducing a perimeter network, referred to as a *DMZ*, that the bastion host resides on. This provides additional security by requiring that an intruder need to bypass two filtering routers in order to gain access to the internal network. Even if the bastion host is compromised, the intruder would still need to get past another packet-filtering router (the internal router) to gain access to the internal network. At best the attacker could gain access to the perimeter network; however, that risk is mitigated by the fact that anything on the perimeter network should be designed as a sacrificial host anyway. While this design provides one of the most secure methods of providing external access to resources, it has some drawbacks, particularly in complexity of design and cost. Figure 2.15 is an example of a screened subnet firewall with a DMZ.

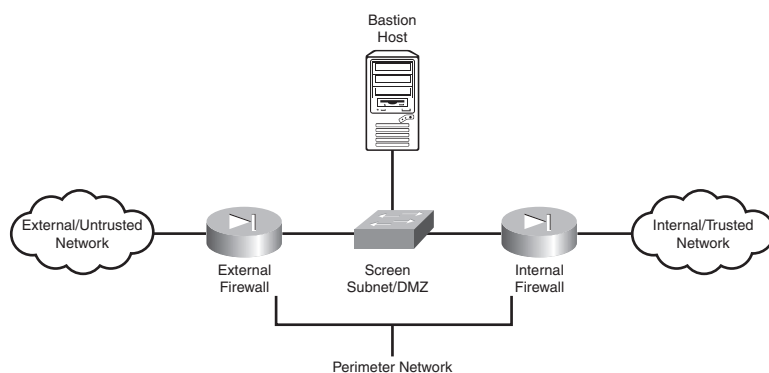


FIGURE 2.15
Screened-subnet firewall (with DMZ).

◆ **Dual homed host firewall**—In a dual homed host firewall system, the bastion host has two interfaces, one connected to the internal network and the other connected to the external network; however, IP forwarding (the ability to route) is disabled. This allows hosts from either network to communicate with the bastion host, but the hosts on the networks cannot communicate with each other via the bastion host. There are a couple of drawbacks to a dual homed firewall system. First, because the bastion host is connected to the internal network, if it is compromised, the intruder would potentially have free run on the internal network. Second, if you decide to allow the bastion host to route, it generally does not perform very well in that role because that is not what it was primarily designed to do. Figure 2.16 is an example of a dual homed host firewall system.




FIGURE 2.16
Dual homed host firewall.

NOTE

The Testing and Verification of Firewall Systems TruSecure maintains an independent firewall testing criteria and a number of excellent FAQs and whitepapers that can provide more detailed information about firewalls. They can be accessed at <http://www.icsalabs.com/html/communities/firewalls/index.shtml>.

Much like firewall types, there are variations and hybrid designs of firewall architectures, but they are all based in part on these for principal designs. An example of this would be a SOCKS server, which is often used to provide proxy based outbound access for clients running SOCKS client software. While this can do a great job of securing access to resources, it has some significant drawbacks in terms of IT support due to the requirement of the SOCKS client on every desktop.

Gateways and Proxies

The term *gateway* has a number of meanings depending on the context used. In some cases a gateway is effectively a router. In other cases, a gateway can mean a device that provides proxy type functionality. In its most basic definition, a gateway provides access to a network or service.

Proxies, on the other hand, provide a very specific function. Proxies are used as intermediary devices between a client and a server, providing the client transparent access to the resources on the server without allowing the client to access those resources directly. As a result, proxies can be used as a security device (for example, an application proxy firewall). Because the traffic between the client and server must go through the proxy, the administrator can restrict and control traffic at a single network location. A common implementation of proxies is to provide outbound Internet access. This allows the administrator to be able to do things like restrict access to sites. Proxy servers will also often cache data, so they can provide better network performance by servicing requests with cached data as opposed to needing to go to the destination for the response.

WAN TECHNOLOGIES

Whereas LAN technologies tend to focus on connecting a large number of systems that are in close proximity to each other to a very fast network, WAN technologies tend to focus on interconnecting LANs and making connections to remote sites and resources. There are three main categories of WAN networks:

- ◆ **Internet**—The Internet is, obviously, the largest WAN on the planet. With roots in the ARPAnet (Advanced Research Projects Agency Network), the Internet is used to provide a global network of resources and access points known as Internet service providers (ISPs). More and more companies are using the Internet as a connection medium, securing the traffic in VPN tunnels.
- ◆ **Intranet**—An intranet is a private network based in concept on the Internet, but it uses company-owned resources for connecting devices and networks. The term *intranet* is also frequently used to refer to the publishing via Web sites of company-specific information. As a result of the security of running on company resources, intranets are generally much more secure than using the Internet.
- ◆ **Extranet**—Similar to an intranet, extranets are used to provide external access to users outside of the company, but they do not allow access from public users. Examples of extranet communications can be company partners who are permitted to access the extranet to gain access only to the information that they need.

There are a number of types of WAN connections to be aware of, as discussed in the following sections.

Dedicated Connections

Dedicated WAN connections exist between two point-to-point sites and generally are available at all times. Once the circuit is paid for, the connection exists around the clock exclusively for the traffic the customer is generating. These connections tend to be *synchronous serial connections*, which simply means that the communication between sites occurs with precision clocking and control bits that specify the beginning and end of transmission characters. The classic example of a synchronous serial connection is a T1 (or E1 in Europe).

Synchronous serial lines are generally available at speeds up to 45Mbps (T3 or E3 speeds). The more common connection speeds and types are

- ◆ **Digital Signal Level 0 (DS-0)**—Defines the framing specification used to transmit data on a single 64Kbps channel over a T1 line.
- ◆ **Digital Signal Level 1 (DS-1)**—Defines the framing specification for transmitting data at 1.544MBps over a T1 or 2.048Mbps on an E1 line.
- ◆ **Digital Signal Level 3 (DS-3)**—Defines the framing specification for transmitting data at 44.736Mbps on a T3 line.
- ◆ **T1**—A T1 carries 24 PCM (Pulse Code Modulations) signals, sometimes called channels, using TDM (Time Division Multiplexing) to achieve a transmission speed of 1.544MBps over a dedicated connection.
- ◆ **T3**—A T3 carries 672 PCM (Pulse Code Modulations) signals, sometimes called channels, using TDM (Time Division Multiplexing) to achieve a transmission speed of 44.736Mbps over a dedicated connection.
- ◆ **E1**—Similar to a T1, E1s are used primarily in Europe and carry data at 2.048Mbps.
- ◆ **E3**—Similar to an E1, E3s are used primarily in Europe and carry data at 34.368Mbps.
- ◆ **OC-x (Optical Carrier X)**—The various optical carriers are a subset of the SONET (Synchronous Optical Network) specification for transmitting digital signals over fiber-optic cable. The base OC rate of OC-1 is 51.84Mbps. The numeric value of the OC rate is multiplied by the base rate to get the speed. OC-3 transmits at 155.52Mbps, OC-12 is 622.08Mbps, OC-24 is 1.244Gbps, OC-48 is 2.488Gbps, OC-192 is ~10Gbps, OC-256 is 13.271Gbps, and OC-768 is ~40Gbps.

These connections are generally considered very secure, because they exist between the two sites and are shared by no one else.

Circuit-Switched Connections

Circuit-switched connections are based on the classic telephone network. When two devices need to communicate between each other, the data network they are using will dynamically bring up the circuits (or connections) that the two devices require in order to exchange data. These circuits are maintained for the duration of the call, which could lead to inefficient use of network resources. For example, if the connection were always left open, it would prevent other connections from being made. Circuit-switched networks tend to use asynchronous serial connections, which simply means that there is no timing of the data stream. Circuit-switched connections tend to use dialup modems and ISDN, and thus are typically used for low bandwidth or backup purposes. Because the connection is established essentially by dialing the destination, provided authentication occurs to allow the connection, circuit-switched is considered a fairly secure connection.

Packet-Switched Connections

Like dedicated connections, packet-switched connections use a synchronous serial method of communications. Where packet switching differs is that the packet-switched network is often shared by multiple systems. The reason for this is simple. Often, a company does not need a dedicated connection between sites with dedicated bandwidth. The cost of maintaining such a connection can be very expensive and by going with packet-switched the company can effectively “time share” the WAN connection. They do this by purchasing a guaranteed amount of bandwidth, for example 128Kbps. Because lots of WAN traffic is small, bursty traffic, the company can have the performance that it needs, but save costs by allowing the underlying circuits to be shared among multiple companies and networks, effectively operating kind of like a party line. No matter what, the company will get the minimum bandwidth it purchased (often times called the CIR (Committed Information Rate), but if more bandwidth is available, the company is able to use it. The classic packet-switched network is frame relay or X.25 with speeds generally ranging from 56Kbps to 2.048Mbps. While not as secure as a dedicated or circuit-switched network, packet-switched networks are still considered a fairly secure WAN medium.

Cell-Switched Connections

Cell-switched networks are similar to packet-switched networks, with one important difference—cell-switched networks are *ATM* (*Asynchronous Transfer Mode*) networks. ATM is a networking standard that uses fixed length 53-byte cells in the transmission of multiple services, such as voice, video, and data. Because of the fixed length cell size, transit delays are reduced because the equipment can be configured to programmatically be prepared for data transmission and receipt. ATM is designed for use on high speed media, for example SONET, T3, and E3 with speed capabilities well into the Gbps capacity. In fact, ATM has no theoretical top speed, but rather relies on the underlying media to establish the rate of transmission. Like packet-switched, ATM is considered a fairly secure WAN technology.

WAN Services

Whereas most LAN connections are based on either Ethernet or Token-Ring, a number of different WAN services provide for WAN connectivity in an internetwork. The following sections discuss these.

Point-to-Point Protocol and Serial Line Internet Protocol

Serial Line Internet Protocol (SLIP) and Point-to-Point Protocol (PPP) are primarily used for providing datalink connectivity over asynchronous (dial-up) and synchronous serial (ISDN or dedicated serial lines such as T1) connections. SLIP was developed first, and provided the ability to authenticate connections before allowing them to be established. A drawback of SLIP is that it only supported IP communications at the Network layer. PPP was developed to replace SLIP, and came with a number of enhancements including multi-protocol support, error detection, and advanced authentication methods.

PPP primarily exists to transport Network layer protocols across a point-to-point connection. Examples of point-to-point connections are dialup, ISDN, and dedicated synchronous serial connections, for example T-1 and T-3 lines. When a device attempts to initiate a PPP connection, three phases of communication occur before data can be transmitted:

- ◆ **Link Establishment Phase**—LCP packets are used to configure and test the link.

- ◆ **Authentication Phase**—CHAP (Challenge Handshake Authentication Protocol), PAP (Password Authentication Protocol), or manual authentication of the connecting devices occurs.
- ◆ **Network Layer Protocol Phase**—PPP uses NCP (Network Control Protocol) to determine what Network layer protocols need to be encapsulated, and are transmitted accordingly.

As mentioned, PPP can use two forms of “automatic” authentication, PAP and CHAP. PAP is the less secure of the two, sending the passwords over the wire in clear text. PAP only performs authentication during the initial connection phase. CHAP on the other hand performs authentication during the initial connection phase, and then periodically revalidates the password for the duration of the connection. CHAP uses an MD5 hash for security of the username and password.

High-Level Data-Link Control

High-Level Data-Link Control (HDLC) is an ISO-based standard for delivering data over synchronous serial lines. HDLC is a bit-oriented datalink protocol that uses frame characters and checksums as part of the data encapsulation. A drawback of HDLC is that there is no authentication. Another problem with HDLC is that, while a standard, it does not provide for specifying the network-layer protocol that was encapsulated. Each vendor developed its own method of identifying the Network layer protocol in use. As a result, while HDLC works great between equipment made by the same vendor, it is often incompatible and thus cannot be used when connecting devices from different vendors. If authentication is required, it is recommended to use PPP instead.

X.25

X.25 is a WAN connection technique that functions at the physical and Data Link layers of the OSI model. X.25 uses virtual circuits for establishing the communications channel between hosts. A very reliable protocol, X.25 has been replaced in many environments by the faster Frame Relay.

Link Access Procedure Balanced

Link Access Procedure Balanced (LAPB) was originally created for use on X.25 networks. LAPB is a bit-oriented protocol, similar to HDLC, and functions by ensuring that frames are correctly ordered and error free.

Frame Relay

Frame relay is one of the most popular WAN connection techniques due to its reliability and support of multiple protocols. Frame relay is based on X.25, but it is considered a faster technology because it leaves error correcting functionality to higher layers. Functioning at the physical and Data Link layer, frame relay provides the communication interface between the DTE (Data Terminal Equipment) and DCE (Data Circuit-Terminating Equipment). Connectivity between two DTEs is provided via the use of virtual circuits, similar to X.25. Frame relay uses DLCIs (Data-Link Connection Identifiers) to identify the end points of communication of a circuit. Frame relay functions at speeds up to 2Mbps and does not use authentication. Like HDLC, if authentication is required it is recommended to use something such as PPP instead.

Synchronous Data-Link Control

Synchronous Data-Link Control (SDLC) is a bit-oriented connection protocol that was designed by IBM for use in mainframe connectivity. SDLC is also used in point-to-point WAN connections. SDLC was largely incorporated into IBM's SNA (Systems Network Architecture) and SAA (Service Application Architecture) for mainframe connectivity and has been largely replaced by HDLC for WAN connectivity.

Integrated Services Data Network

Integrated Services Data Network (ISDN) was developed as a standard for transmitting digital signals over standard telephone wires. ISDN functions at basic rate interface (BRI) speeds up to 128Kbps and primary rate interface (PRI) speeds up to 1.544Mbps. Two levels of service are defined by ISDN:

- ◆ **BRI**—BRI is intended for small office and home user usage. BRI supports the use of one 16 Kbps D channel (Delta channel), which is used for carrying signaling and control information, and two 64Kbps B channels (Bearer channels), which are used to transmit voice, video, and data.
- ◆ **PRI**—PRI is intended for greater usage, and for connecting multiple BRI connections. PRI uses a single 64Kbps D channel and 23Mbps B channels (in Europe there are 30 B channels).

ISDN is typically used in conjunction with PPP which allows for both B channels to be bonded together in a multilink connection, providing for 128Kbps throughput (the sum of both B channels). Still a viable backup WAN routing connection and primary low speed WAN connection used in small office environments, ISDN has been pushed to the wayside of the home market with the advent of broadband technologies like DSL (Digital Subscriber Line) and cable modem.

Digital Subscriber Line

Digital Subscriber Line (xDSL) is a relatively new technology that supports the broadband transmission of data at high speeds, currently up to about 53Mbps, over the existing telephone network. xDSL is rapidly becoming the standard for inexpensive remote connectivity, particularly for home users and telecommuters. There are four primary types of DSL:

- ◆ **Asymmetric Digital Subscriber Line (ADSL)**—ADSL is designed to deliver higher download speeds, from 1.5 to 9Mbps, with upload speeds ranging from 16 to 640Kbps. ADSL is supported up to distances of 18,000 feet from the central office using a single line.
- ◆ **Single-line Digital Subscriber Line (SDSL)**—SDSL is designed to provide downstream and upstream speeds of 1.544Mbps. The practical distance limitation of SDSL is about 10,000 feet from the central office using a single line.
- ◆ **High-rate Digital Subscriber Line (HDSL)**—HDSL also functions at speeds of 1.544Mbps, but HDSL uses two lines allowing it to function in a full duplex mode. HDSL is often used by providers for actually providing T1 connectivity.

HDSL is able to run at distances up to 12,000 feet from the central office.

- ◆ **Very-high Digital Subscriber Line (VDSL)**—VDSL is designed to deliver network speeds of 13 to 52Mbps downstream and 1.5 to 2.3Mbps of upstream over a single wire. Unfortunately, the operating range of VDSL is only 1000–4500 feet from the central office.

Switched Multimegabit Data Service

Switched Multimegabit Data Service (SMDS) is a high-speed packet-switching technology for use over public networks. It is provided for companies that need to send and receive large amounts of data on a bursty basis, providing for connectionless communications. It is a bandwidth-on-demand technology.

High Speed Serial Interface

High Speed Serial Interface (HSSI), sometimes called “hissy,” provides for an extremely fast point-to-point connection between devices, but the distance limitation is no more than 50 feet. HSSI can transmit data at speeds of 53Mbps, allowing it to be used to connect devices at T3 or OC-1 speeds. HSSI is often used to interconnect LAN equipment for backup and fault tolerant network uses.

WAN Devices

Now that you have seen the theory and architecture that WAN connections are built with, as well as the physical interconnection methods and networking types, you need to look at the components and technologies that enable WAN connectivity. These are

- ◆ **Routers**—Although routers are a LAN device, they are also used extensively on WANs to provide routing between subnets.
- ◆ **WAN switches**—WAN switches operate at the Data Link layer of the OSI model, but that is where their similarity with LAN switches ends. Typically used on the carrier networks, WAN switches connect private data over public circuits.

- ◆ **Multiplexors**—Called MUX for short, a multiplexor enables more than one signal to be transmitted simultaneously over a single circuit.
- ◆ **Access servers**—Access servers are often used for dial-in and dial-out access to the network. We will look at remote access in more detail in a moment.
- ◆ **Modems**—A modem is responsible for converting digital and analog signals, allowing digital data to be transmitted over analog phone lines.
- ◆ **CSU/DSU (Channel Service Unit/Data Service Unit)**—CSU/DSUs are digital interface devices that are used to terminate the physical connection on a DTE device (for example, a router) to the DCE (for example, a WAN switch).

PROVIDING REMOTE ACCESS CAPABILITIES

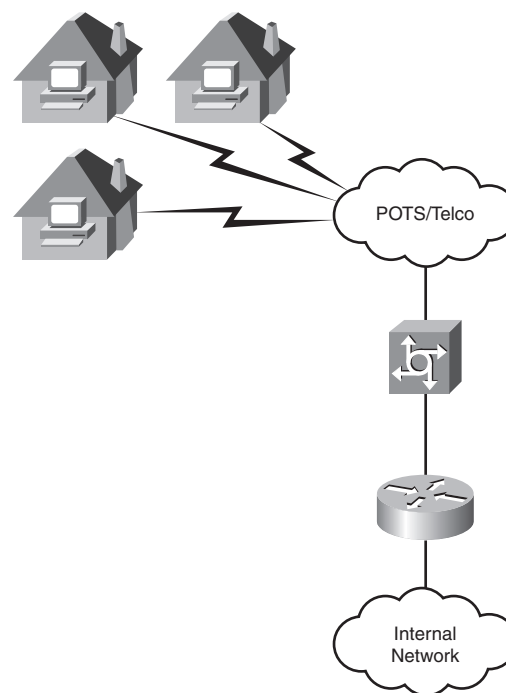
One of the most dangerous items of network design is the increasing need for remote-access capabilities for workers. With the advent of telecommuting, the strain of providing secure networks has become even more difficult to manage. I recently read a trade article that mentioned that 25% of IBM's global workforce telecommutes, and a sizable number of those users do not even have a formal desk. Let's look at a few remote access techniques and technologies.

Client-Based Dial-in Remote Access

Client-based dial-in remote access, or dial-up access, is the classic remote access scenario. Users work from home (telecommuting) or on the road and need access to corporate resources such as email and databases. Typically the client will run some sort of access software on a PC and connect to the corporate network via a hardware device or server. One method of connectivity is to dial in to the corporate network via a modem, thus providing connectivity to the corporate network. This requires that the company maintain some sort of modem bank that their users can call.

A method of remote access that is becoming more and more used is dialing into an ISP (Internet service provider) via the POTS (plain old telephone system) or local TELCO (telecommunications company) and creating a VPN (virtual private network) tunnel across the Internet to a VPN server on the corporate network. We will look at client-based VPNs in a moment. Figure 2.17 illustrates how a client-based dial-in connection is made.

FIGURE 2.17
Client-based dial-in connection.



Using Tunneling As a Security Method

Tunneling is the process of transmitting one protocol encapsulated within another protocol. This allows for the transmission of data that might not be supported on the network via the data that is supported. Tunneling is often used to create a secure channel (a VPN) over an otherwise insecure network, typically the Internet. Tunnels usually designate two endpoints of communications, and then encapsulate the data to be transmitted within some other packet format, thus creating the tunnel from point-to-point.

An important thing to understand about tunneling is that it does not replace encryption/decryption of the data. Although all good tunneling implementations should have some form of encryption built into the tunneling mechanism, the data will still be accessible without built-in encryption.

Numerous tunneling techniques are used today, but most share the common goal of being used to provide VPN connections. Point-to-Point Tunneling Protocol (PPTP) is a tunneling technique that is very popular due to the support and development of Microsoft, and the native inclusion of it on many Microsoft operating systems. PPTP is typically used to create connections across the Internet between devices. PPTP provides for data encryption capabilities. Cisco also has a popular tunneling technique using GRE (Generic Routing Encapsulation) which is typically used for providing VPN connections as well. IPSec (Internet protocol security) is often used in conjunction with GRE to provide for data encryption.

Virtual Private Networks

A virtual private network (VPN) is simply the use of a “tunnel,” or secure channel, across the Internet or other public network. The data within the tunnel is encrypted, thus providing security and integrity of the data against outside users. When implemented properly, VPNs can provide a cost-effective method of providing secure remote office, small office, and remote user connectivity. I like to think of a VPN as an armored car. The money (data) is encapsulated in an armored car (secure packet format) so that it can be transmitted over the public streets (Internet) with a relatively low likelihood of an unauthorized person gaining access to it.

VPNs exist in one of two forms: client-based and site-to-site.

Client-Based VPNs

Client-based VPNs provide remote access to users. Users runs some form of VPN client software on their computers, which allows them to connect to the corporate network as if they were a node on that network. Unlike site-to-site connections, client-based VPNs rarely allow for systems other than the one running the client software to connect with the VPN.

The remote client becomes a *virtual node* of the network to which it is connecting. This has become a much more popular method of connecting than client-based dial-in connections, for two reasons: First, most users have Internet access already, particularly those with broadband access; second, by using an ISP, the company can avoid long distance charges by having its users call a local number and then use the Internet for the connectivity to the corporate network.

Site-to-Site VPNs

Site-to-site remote access connections have come into use as a mechanism for connecting remote sites via the Internet. A site-to-site VPN is a permanent or semi-permanent connection between two devices, typically firewalls or routers. Site-to-site connections link up remote offices across the Internet. Computers on the remote LAN require no special software to communicate with the network to which the VPN connects. Rather than paying for an expensive site-to-site or packet-switched WAN connection for remote access, companies have begun using the Internet as their WAN connection with a VPN used to secure the traffic. Although this can provide a relatively cheap method of connectivity for small remote sites and home offices, particularly using high-speed broadband technologies such as DSL (Digital Subscriber Line) and cable modem, you must remember that the Internet is not a reliable connection. If a reliable site-to-site remote access connection is required, you really need to go with a traditional WAN solution. The benefit of the site-to-site remote access solution is that individual clients on the remote network require no special software or configuration to provide remote access capabilities. Typically a router or VPN hardware device handles all client requests, forwarding them to the Internet or to the remote site as required. This is known as *split tunneling*. Figure 2.18 illustrates how client-based and host-to-host VPN connections would be connected.

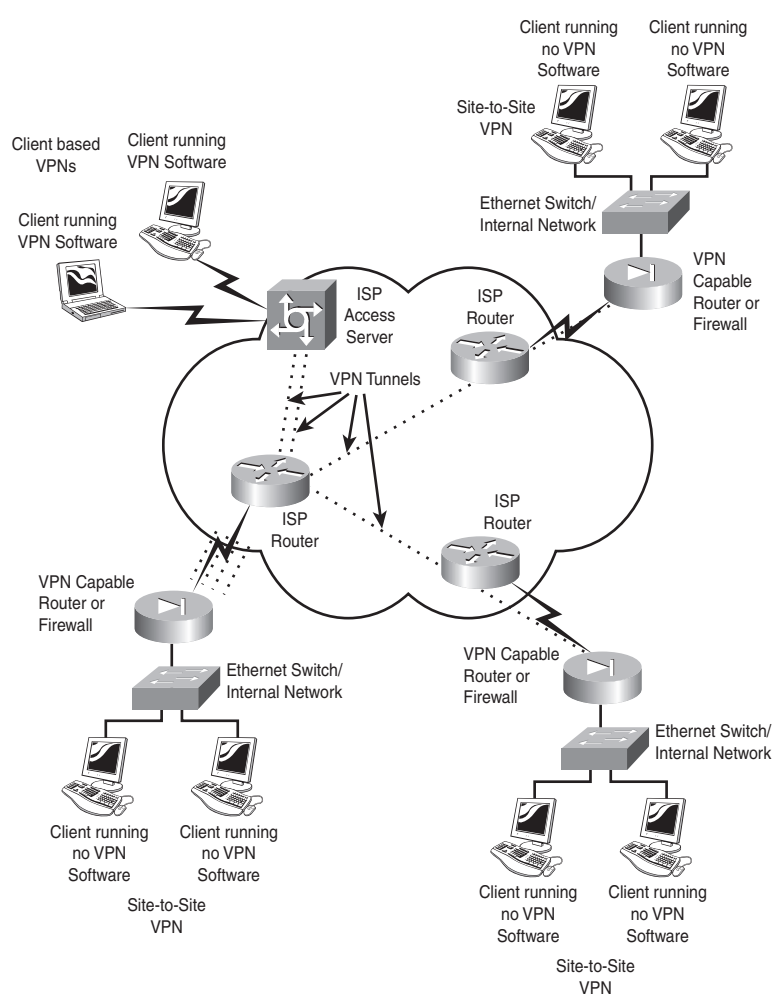


FIGURE 2.18
VPN connections.

VPN Protocols

Three primary technologies are used for providing remote access VPN capabilities:

- ◆ **PPTP**—PPTP is a Microsoft-developed technology that provides remote access by encapsulating PPP inside a PPTP packet. PPTP uses the PPP authentication mechanisms of PAP, CHAP, or MS-CHAP for authentication and RSA RC4 and 40-bit or 128-bit session keys and encryption. PPTP supports multi-protocol tunneling.

- ◆ **L2TP (Layer 2 Tunneling Protocol)**—L2TP is similar in function to PPTP, but it does not use any vendor-specific encryption technologies. In addition, L2TP supports the use of RADIUS (Remote Authentication Dial-In User Server) and TACACS (Terminal Access Controller Access Control Service) for authentication, and IPSec (Internet Protocol Security) and IKE (Internet Key Exchange) for encryption and key exchange respectively. L2TP supports multi-protocol tunneling.
- ◆ **IPSec**—IPSec is a network-layer encryption and security mechanism that can be used as a standalone VPN solution, or as a component of an L2TP VPN solution. IPSec supports the use of DES (Data Encryption Standard) and 3DES (Triple DES), although because the DES scheme was successfully hacked in 1999, it is highly recommended that you only use 3DES. The integrity of the data can be provided via 128-bit MD5-HMAC (Message Digest 5—Hash Message Authentication Code) or 160-bit SHA-HMAC (Secure Hash Algorithm—Hash Message Authentication Code). IPSec supports the use of AH (Authentication Header) security, in which the IP header is secured but the data is not, or ESP (Encapsulation Security Payload) in which the entire packet is encrypted and secured.

Remote Access Authentication

RADIUS is a UDP-based de facto industry standard for providing remote access authentication via a client/server model. When the client attempts to connect to the network, it is prompted for a username and password that is checked against a user database existing on a network server. RADIUS uses a combined authentication and authorization profile, which means that RADIUS access is typically “all or none.” You are either allowed to connect, or you are not.

TACACS (Terminal Access Controller Access Control System) is an older authentication technology that has been largely marked “end-of-life,” which means that people should refrain from using it.

TACACS+, which sounds similar, is actually an entirely new protocol. Similar in function to RADIUS, TACACS+ differentiates itself from RADIUS by separating the authentication and authorization capabilities, as well as using TCP for connectivity. As a result, TACACS+ is generally regarded as being more reliable than RADIUS.

NETWORKING PROTOCOLS

Protocols are simply the rules by which something functions. In the case of network protocols, these are the rules that control how data is processed. Protocols often have OSI layer-specific functionality that they are responsible for. As discussed in the following sections, there are a number of protocols that a security professional should be aware of in network environments.

Transmission Control Protocol/Internet Protocol

Transmission Control Protocol/Internet Protocol (TCP/IP) is the foundation on which virtually all networking today occurs. TCP/IP is actually a suite of protocols that was developed by the Department of Defense to provide a highly reliable and fault tolerant communications infrastructure. TCP/IP was designed following a four-layer architectural model, as opposed to the OSI seven-layer model as illustrated in Figure 2.19.

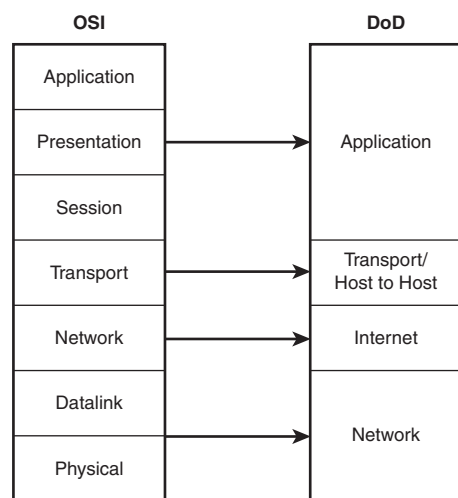


FIGURE 2.19

The DoD model versus the OSI model.

The four layers of the DoD model are as follows:

- ◆ **Application layer**—The Application layer loosely maps to the top three layers of the OSI model, and provides for the applications, services, and processes that run on a network.

- ◆ **Transport layer**—Sometimes referred to as the host-to-host layer, the Transport layer is responsible for handling the end-to-end data delivery on the network. It loosely maps to the Transport layer of OSI.
- ◆ **Internet layer**—The Internet layer maps loosely to the Network layer of the OSI model and provides logical addressing and routing of IP datagrams on the network.
- ◆ **Network layer**—The Network layer maps loosely to the datalink and Physical layers of the OSI model. The Network layer is primarily responsible for the physical delivery of data on the network.

Let's look at the four layers in more detail.

Application Layer Protocols

There are a number of different application-layer protocols as services. They are largely responsible for providing user access to the network. Some of the more common protocols are

- ◆ **Bootstrap Protocol (BootP)**—BootP is used to provide for automatic configuration of diskless workstations by looking up the client MAC address in the BootP file. When it finds the entry, it sends the client the necessary information needed to complete the system boot process.
- ◆ **File Transfer Protocol (FTP)**—FTP is used to send and receive files between two systems. FTP provides for authentication, albeit using clear-text passwords, and does not provide for the remote execution of programs.
- ◆ **Line Printer Daemon (LPD)**—LPD, when used in conjunction with LPR (Line Printer Remote), is used for connecting to network-attached print devices.
- ◆ **Network File System (NFS)**—NFS is a file-sharing protocol, typically used in Unix environments.
- ◆ **Post Office Protocol 3 (POP3)**—POP3 provides for the connecting to and receipt of email from a mail server to the email client.
- ◆ **Simple Mail Transfer Protocol (SMTP)**—SMTP provides for the delivery of email across servers throughout a network.

Whereas POP3 is primarily responsible for the receipt of email, SMTP is primarily responsible for sending email.

- ◆ **Simple Network Management Protocol (SNMP)**—SNMP is designed to support the transmission and collection of management information and statistics for network devices. SNMP can be configured to notify when a network event occurs by sending *traps*. SNMP also provides mechanisms to allow network administrators to make changes on remote systems via *set* operations. The information that a device can report on or change is maintained via files known as MIBs (Management Information Bases), which are databases containing the information that SNMP is aware of.
- ◆ **Telnet**—Telnet provides remote command-line functionality across an IP internetwork. Telnet is a terminal-emulation program that can be used to remotely execute commands and run applications, but it cannot be used for file transfers.
- ◆ **Trivial File Transfer Protocol (TFTP)**—A subset of FTP, TFTP is used to provide file-transfer services. TFTP lacks FTP's more robust features such as authentication and directory browsing. TFTP is commonly used to update router and switch configurations and software, but it is inherently insecure and should only be used with caution.
- ◆ **X Window**—X Window is a protocol that facilitates the remote display of the GUI, primarily in a Unix environment.

Transport Layer Protocols

A number of protocols reside at the Transport layer, the most significant being TCP (Transmission Control Protocol) and UDP (User Datagram Protocol). For reliable, connection-oriented data services, TCP/IP uses TCP. For unreliable, connectionless data services, TCP/IP uses UDP. Because TCP requires connection establishment, it is considered more reliable and secure than UDP. Following are detailed descriptions of TCP and UDP:

- ◆ **Transport Control Protocol (TCP)**—Primarily responsible for creating connection-oriented, reliable end-to-end communications between host systems. TCP does this via a series of synchronizations (called *SYNs*) and acknowledgements (called *ACKs*)

prior to data transfer. This is sometimes called a TCP three-way handshake. As a part of the handshake process, TCP also establishes a method of periodically checking to ensure that the data is being reliably delivered via a mechanism known as *windowing*. When the two host systems decide to communicate, they establish an amount of data that can be sent before an acknowledgement must be received. In doing so, the source host can ensure that the destination host received the data properly. If the source host does not receive an acknowledgement within a predetermined amount of time, the source will assume that the data was lost and retransmit it. Because the data is broken into segments, and the segments may arrive at the destination out of order, TCP uses sequence numbers so that the destination knows in what order the segments should be put back together. TCP also manages the flow of data to reduce congestion, overloading, and loss of packets. TCP is defined in RFC 793 and is updated by RFC 3168.

- ◆ **User Datagram Protocol (UDP)**—Primarily responsible for connectionless, unreliable end-to-end communications between systems. The first thing that people always ask about UDP is, “Why would someone want to use an unreliable protocol?” The answer is really quite simple—when the receipt of the data is not that important, or when the overhead of ensuring the reliable delivery of data is too high. For example, if you have a communications method that uses frequent small transactions (for example, many network management applications), the overhead of establishing, maintaining, and tearing down the session every time the two hosts communicate can actually be more data than the actual network management data. In other cases—for example, NFS—reliability is guaranteed by higher layer protocols, so using TCP would be excess overhead. A relatively new type of application that also frequently finds itself suited for UDP is streaming audio and video, because you don’t want to send previously lost packets. If you did, the sound and video could arrive out of order. UDP is defined in RFC 768.

TCP and UDP use port numbers as the endpoints of communications to define the upper-layer applications and conversations that are occurring. The list of all registered TCP and UDP port numbers can be located at <http://www.iana.org/assignments/port-numbers>.

Reviewing TCP and UDP

REVIEW BREAK

TCP and UDP provide the mechanism that hosts use to transport data between hosts across the network. Table 2.2 compares TCP and UDP, highlighting the key functions of each.

TABLE 2.2
COMPARISON OF TCP AND UDP

<i>TCP</i>	<i>UDP</i>
Acknowledged data transfer	Unacknowledged data transfer
Uses sequencing	Does not use sequencing
Connection-oriented	Connectionless
Reliable	Unreliable
Higher overhead	Lower overhead

Internet Layer Protocols

The Internet layer *is* TCP/IP. Virtually every other protocol used in networking is designed to specifically interface to and support the various Internet-layer protocols. Some of the more common Internet-layer protocols are

- ◆ **Internet Protocol (IP)**—As mentioned, IP is responsible for handling the logical addressing of hosts with the use of IP addresses. IP addresses consist of 32 bits of data, separated into four 8-bit sections known as octets. IP is considered an unreliable delivery mechanism, which is fine because TCP can provide reliability if desired. With IP, data can be delivered never, once, many times, in order, or out of order, and IP does not care.
- ◆ **Internet Control Message Protocol (ICMP)**—ICMP is a management and control protocol for IP. ICMP is responsible for delivering messages between hosts regarding the health of the network. This information could be reachability of hosts as well as routing information and updates. Many IP diagnostic tools use ICMP, such as PING (Packet Inter-Network Groper) and Traceroute. ICMP is defined in RFC 792 and is updated by RFC 950.

- ◆ **Address Resolution Protocol (ARP)**—All hosts require the physical and logical address of the host with which they want to communicate. Because the IP address is known by the source, but the hardware address may be unknown, ARP is used to discover and maintain a list of IP addresses and their respective MAC addresses. ARP functions by sending a broadcast, known as an ARP request, to the entire subnet to discover the MAC address of the known IP address. The host that owns the IP address in question will respond with its MAC address, thereby allowing for communications between hosts.
- ◆ **Reverse Address Resolution Protocol (RARP)**—As the name would imply, RARP performs the exact opposite function of ARP. Sometimes the MAC address is known, and the IP address is what needs to be determined. RARP is commonly used in diskless workstations, where the system knows its MAC address but needs to get the IP configuration information. In these cases, a RARP server can respond with the required information.

PROTECTING THE INTEGRITY, AVAILABILITY, AND CONFIDENTIALITY OF NETWORK DATA

Now that we understand how networks are built and the technologies and devices that run them, we are ready to look at some of the technologies and techniques that we can use to protect the integrity, availability, and confidentiality of transactions over networks.

The CIA Triad

Information systems security can be addressed by applying the concepts of the CIA triad in the following ways:

- ◆ **Confidentiality**—Confidentiality is simply ensuring that the data transmitted is only able to be read by the intended recipient.

Chapter 2 TELECOMMUNICATIONS AND NETWORK SECURITY 131

The loss of confidentiality can occur in many ways—for example, through the intentional disclosure of information or through lax security procedures. Confidentiality of data can be protected by employing some of the following techniques:

- Network security protocols
 - Network authentication services
 - Data encryption services
- ◆ **Integrity**—Integrity of data is simply the assurance that the data that was received is the data that was transmitted. The data should not be altered, and if it was, there needs to be some method to identify that the alteration occurred. A number of techniques can ensure data integrity:
- Nonrepudiation of message source
 - Firewall systems
 - Communications security
 - Intrusion detection systems
- ◆ **Availability**—Availability is a concept that can be applied to create reliability and stability of network systems and applications. Availability ensures that data is available when required. Although availability is not traditionally considered an aspect of the security professional's area of responsibility, with the prevalence of Denial of Service attacks, the need for data to be “always on” is critical. Some techniques of assuring availability are
- Fault tolerance of disks, systems, and backups
 - Acceptable log-in and process performance
 - Reliable and functional security processes and mechanisms

For more information on the CIA Triad, see Chapter 6, “Security Architecture and Models.”

Security Boundaries and Translating Security Policy to Controls

One of the most effective ways to handle security is to identify needs and risks, and define boundaries that separate services from potential harm. Most networks can be defined with three major groupings:

- ◆ **External subnets**—External subnets contain those resources that the security administrator has no control over. Systems placed on or directly connected to external subnets (that is, the Internet) must be hardened and built from the perspective that they will be under a constant assault from malicious users. These systems should run the bare minimum services and applications required to perform a task.
- ◆ **Internal subnets**—Internal subnets contain those resources that the security has control over. Unfortunately, most people treat internal subnet security a distant second to external security, even though most hack attempts occur on internal networks. The key to securing internal subnets is the separation of resources (that is, place HR data on a server that only HR can access), the auditing of transactions (run network-based IDS and packet sniffers), and the definition of an enforceable security policy.
- ◆ **Screened subnets**—Screened subnets, sometimes referred to as DMZs, are used to provide limited access to external users while still maintaining some degree of control over the resources. An example would be allowing external access to a server via port 80 (HTTP) but preventing all other external access via packet filtering.

In addition to separating network boundaries, it is also prudent to define groupings of processes into domains and types based on least privilege. This is known as *type enforcement*. You should group resources based on how the resource can be used and by whom. Access should then be granted only to those users who need the data, and even then the users should gain only the minimum access required. Once you have determined the groupings of resources, you can define how the resources should be distributed. In some cases it may be necessary to physically separate the resources onto different servers and subnets to provide for granular audit and access control.

In other cases, it may be suitable to use a single server with file permissions and security to prevent unauthorized access. Many of these principles are based on what is known as the *Rainbow Series* of books.

Trusted Network Interpretation

The Department of Defense developed a series of books known as the *Rainbow Series* due to the various colors of their covers. The cornerstone is a book known as the *Orange Book*, which defines the TCSEC (Trusted Computer Security Evaluation Criteria). The other books, including the *Red Book*, expound upon the concepts introduced in the orange book. Although some of the concepts are a little dated, many of the fundamental principles are still very applicable to today's technologies.

According to TSCEC, system security is defined by four broad classifications:

- ◆ **Division D**—Specifies the minimal protection is available, or that the system has failed to meet all other classification.
- ◆ **Division C**—Specifies that, through the use of auditing, discretionary protection and accountability of subjects and the actions they initiate are covered. There are two subclasses of division C:
 - **C1**—Systems at this level satisfy discretionary security by providing for the separation of users and data.
 - **C2**—Systems at this level provide a more granular degree of access through the use of login procedures, auditing of security events, and resource isolation.
- ◆ **Division B**—Specifies that mandatory access control rules are required. Systems in this division are required to carry sensitivity labels with major data structures in the system. Division B has three subclasses:
 - **B1**—B1 systems require all of the features of a C2 system. In addition, an informal statement of the security policy model, data labeling, and mandatory access control over named subjects and objects must be present.

- **B2**—B2 systems require a formal, structured security policy model that requires the discretionary and mandatory access control be extended to all subjects and objects in the system.
 - **B3**—B3 systems require the use of security domains to mediate all accesses of subjects to objects to ensure tamperproof function.
- ◆ **Division A**—These systems use formal security verification to assure that all of the security controls employed can effectively protect classified or other sensitive information via a stringent design verification. Division A has one subclass:
- **A1**—Functionally no different than B3, A1 specifies a formal design specifications and verification techniques are used, resulting in a high degree of security.

The key to these security boundaries and practices is the establishment of an effective, enforceable security policy. Unfortunately, a security policy is usually defined at some point in the future, even though in practice it should exist prior to any security implementations to ensure for the structured design of a security practice. The security policy should clearly define what is and is not permitted by both users and administrators. In addition, the security policy should serve as the guideline for defining the types of resources and access that users require to those resources. The security policy should define the procedures that should be followed in the event of a compromise. The last thing that a security professional needs is to be “making the rules as they go” in the event of a compromise. When everything goes wrong on the network, the security policy should serve as a point of reference and guidance in what is surely a hectic time.

Ultimately, however, all the security preparation in the world is pointless unless management buys into and enforces security procedures. I know of a company that had an exhaustive virus protection policy, but its R&D department did not adhere to it because the virus protection software could potentially invalidate a test environment. The company was hit by the CodeRed and Nimda viruses, which, combined, caused a loss of approximately a week of business work. The company literally shut its doors and no work could be performed until the virus outbreak was under control and all systems were cleaned and patched accordingly. Quite surprisingly, only a week later a virus policy came out that was enforceable on all systems, including R&D systems.

Unfortunately, many times, people are unwilling to provide the necessary enforcement of policies until a very painful situation has occurred.

Network Layer Security Protocols

Traditionally, encryption occurs at the Presentation layer. In an effort to optimize and speed up encryption and decryption functions, numerous protocols have been designed that function at lower layers of the OSI model. Some of them are

- ◆ **IPSec**—IPSec is the most predominant Network layer encryption protocol in use today. Based in part on SWIPE, IPSec provides two choices of security: AH (Authentication Header), in which the sender is authenticated but the payload is not encrypted; and ESP (Encapsulated Security Payload), in which not only is the sender authenticated, but the data payload is encrypted. As a result, ESP is considered the more secure of the two. Key management of IPSec is often handled by the ISAKMP/Oakley (Internet Security Association & Key Management Protocol) protocol. IPSec also functions in two modes, tunnel and transport. Tunnel mode is used to encapsulate the entire original IP datagram for use in situations where the protected datagrams are sourced or destined to systems that do not use IPSec—for example, in the case of a VPN. Transport mode encapsulates the upper layer (Transport layer and above) data of the original packet and is used in cases where the end points of communication both support IPSec—for example, a client connecting to a server.

A drawback of IPSec is that it is largely incompatible with NAT (Network Address Translation). IPSec requires that data integrity not be compromised, and NAT by design translates data midstream between hosts; because of this, when the destination system cannot validate the integrity of the traffic due to the source address changing, the data is therefore dropped. Some vendors work around this by encapsulating IPSec traffic in TCP or UDP in order to pass it through a NAT device. If that sounds a little confusing, it is. The original IP datagram is encapsulated in an IPSec datagram, which is encapsulated in a TCP datagram, which is finally able to be transmitted.

The whole thing reminds me of the Russian nesting dolls where you open one to find another, open it only to find another, open it only to find another. See Chapter 6 for more information on IPSec.

- ◆ **SWIPE**—SWIPE is a predecessor to IPSec. SWIPE provides encryption at the Network layer by encapsulating the original packet within the SWIPE packet. SWIPE does not have policy or key management functionality built into the protocol.
- ◆ **Simple Key Management for Internet Protocol (SKIP)**—SKIP is a stateless Network layer encryption mechanism developed and used primarily for SUN Solaris environments, though it functions on Windows-based systems as well. SKIP is able to encrypt data without needing a prior message exchange between hosts in order to establish a secure channel. Consequently, SKIP can be used in simple communications environments.

Transport Layer Security Protocols

Secure Socket Layer (SSL) is an open, non-proprietary Transport layer encryption method that is supported by both firewalls and tunneling. SSL provides for data encryption, server authentication, data integrity, and optional client authentication via TCP/IP. SSL is used primarily for HTTP traffic and securing the communications between Web browsers and Web servers. SSL uses digital certificates for server authentication, encryption for transmission privacy, and end-to-end connections to ensure the data integrity.

The successor to SSL as a transport-layer security protocol is likely going to be handled by the TLS (Transport Layer Security) protocol. Although built on SSL 3.0, TLS does not support SSL directly. TLS is defined in RFC2246, currently a proposed standard by the IETF (Internet Engineering Task Force).

Application Layer Security Protocols

A number of protocols exist at the Application layer for the purpose of securing specific applications. Because of the nature of email transmissions, a few protocols have been developed expressly for the purpose of securing email transmissions including:

- ◆ **Secure/Multipurpose Internet Mail Extensions (S/MIME)**—As the name would imply, S/MIME is a specification for securing email transmissions. Based on MIME and using RSA encryption, S/MIME provides for cryptographic security through MIME encapsulation of digitally signed and encrypted objects. S/MIME ensures that authentication, non-repudiation, message integrity, and confidentiality occur.
- ◆ **Privacy Enhanced Mail (PEM)**—PEM is defined in RFCs 1421, 1422, 1423, and 1424. PEM provides for message encryption and authentication by using symmetric (secret-key) and asymmetric (public-key) encryption methods for encryption of data encryption keys. Unfortunately, because PEM uses a proprietary form of RSA encryption, it is rarely used.

In addition to email, you also have the increasing use of online purchasing and bill payment. To protect these transactions, Secure Electronic Transaction (SET)—a framework for protection against fraud—was developed to provide a framework for protecting the use of credit cards used in Internet transactions against fraud. SET uses a subset of a PKI (Public Key Infrastructure) to provide for the confidentiality and integrity of the cardholder data, while at the same time providing for the authentication of the card.

Network Monitoring and Packet Sniffers

As we have seen, all devices on a segment potentially receive every signal, but they discard the packets that are not addressed to them (either specifically or via broadcast and multicast addressing). Because of this fundamental nature of networking, though, a user can run software known as network-monitoring or packet-sniffing software and capture the data on the segment. This allows a user to see all data on the segment, even if that data is not destined for the system that the network monitoring software is located on. This can be a huge security risk, as one would expect, but it can also be an excellent troubleshooting tool.

I like to think of a packet sniffer, when used as a troubleshooting tool, as a translator, translating the language that computers use (for example, TCP/IP) into something that I can understand.

With a packet sniffer, the security admin can see the exact format of frames and packets, which can be useful if you want to block a certain type of packet format. For example, peer-to-peer file-sharing utilities—such as Morpheus, KaZaa, and Napster—are a current bane to many network admins. The problem is that many of these programs use what would normally be opened application ports for their communications mechanism (for example, using TCP port 80, which is typically used for HTTP communications). However, the data packets have a unique format that does not necessarily match what normal Web browsing traffic looks like. By running a packet sniffer and observing the traffic patterns that the software uses, the security administrator can then configure perimeter security devices such as firewalls and proxies to drop the specific frames that match the pattern of the prohibited software. This is sometimes known as *pattern-based application recognition*, or what Cisco calls *Content-Based Access Control (CBAC)*.

When used by a malicious user, a packet sniffer can provide information that the user would not otherwise have been able to gain access to. For example, when I lecture on how to use a packet sniffer, I will often capture data for about 10 minutes during class. Normally, at least one of the students is checking email, typically over an unencrypted Web site. I will then find the packets that were part of the user's connection to the Web site and will often read back part of the email message (with the victim's permission, of course) to demonstrate how easy it is for a malicious user to gain access to privileged information. In another case, I was working on a product that displayed information via a Web interface. The main Web page used authentication to validate the user, but the data Web pages did not. The data Web pages were used to export the data to another software program. When I explained that this was a security hole, because someone could connect to the data URLs unchallenged, I was told that because the URLs contained a unique object identifier, it would be virtually impossible to guess the URL. I bet lunch that I could gain access to this information in no more than 10 minutes, without needing to guess. I set up a packet sniffer and then proceeded to start the software that the data is exported to. Because it collected data on a 10-minute cycle, I needed to only wait for the packet sniffer to capture the URL, object identifier included, and thus earn a rather nice steak lunch.

One thing to be aware of when using a packet sniffer is how they function in a switched environment. If you recall, switches optimize traffic by making each port on a switch effectively its own segment. As a result, data signals are only sent to the ports on which the destination system is located. Using a packet sniffer to capture data belonging to hosts other than the host running the packet sniffer requires that the host be able to receive the signal containing the data. By default a switch is not going to allow this to happen. To get around this, most switches are configured with a feature called *port spanning* or *port mirroring*. Here is how it works. Let's say that you want to capture the data that is coming to or from a system connected to port 1 on the switch. The system running the packet-sniffing software is connected to port 10. You can configure the switch to forward data that is on port 1 to port 10, thereby allowing the packet sniffer to capture the data that is going to or from the host connected to port 1. It's important to be aware that the NIC with which you are going to be monitoring the traffic will need to be fast enough to keep up with the quantity of traffic that it might receive. In general, it is recommended to use a dedicated NIC for monitoring traffic, and have a completely different NIC for actually communicating with the sniffer.

Intrusion Detection

Intrusion detection is the process of monitoring systems for evidence of an intrusion or misuse. You accomplish this by collecting information from numerous sources and then analyzing the information for symptoms of a security compromise. This information can then be used to alert administrators to determine the relevance and severity of the incident. It is important to note that intrusion detection is not intrusion response. Intrusion response occurs after the event has been properly detected. Intrusion Detection Systems (IDSs) are responsible for performing the following tasks:

- ◆ Monitoring and analyzing user, system, and network access
- ◆ Auditing system configurations and vulnerabilities
- ◆ Assessing the integrity of system and data files
- ◆ Recognizing activity patterns that would seem to indicate an incident

- ◆ Analyzing abnormal use patterns
- ◆ Operating system auditing

In addition, some advanced IDSs can provide

- ◆ Automatic patching of vulnerable systems through recovery actions and scripting
- ◆ Installing and monitoring decoy servers to gather information

Security professionals should be aware of two fundamental variations of IDS: network- versus host-based IDSs; and knowledge- versus behavior-based IDSs.

Network-based IDSs are essentially raw packet-parsing engines, basically a network sniffer on steroids. Network-based IDS compares captured traffic against some known database or pattern of attacks to ascertain whether a potential situation is occurring. Network-based IDS is typically deployed to monitor traffic on network segments—for example, backbone and perimeter network segments. They capture traffic in promiscuous mode, allowing it to capture all traffic on the segment, and will generally analyze the packets in what is considered real time.

Host-based IDS is a little more complex to define because these systems are often system-centric in their design. Most host-based IDS are designed to monitor logins and processes, typically through the use of auditing system logs. Host-based IDS is designed to specifically identify inappropriate activity on the host system only. They are typically agent-based, which means that an agent is required to be running on the system that is being monitored. As a result, host-based IDS can be difficult to deploy and manage.

Knowledge-based IDS can be network- or host-based. A knowledge-based IDS maintains a database of known attacks and vulnerabilities (in other words, knowledge) and detects whether attempts to exploit these vulnerabilities are occurring. Knowledge-based IDS is more common than behavior-based IDS and are sometimes referred to as signature based. Some benefits of knowledge-based IDS are

- ◆ Low degree of false positives
- ◆ Alarms are standard and easy to understand, because they are based on known attacks and exploits

Some drawbacks are

- ◆ Resource intensive, as the knowledge-based IDS must be constantly updated to detect new exploits
- ◆ New attacks can go unnoticed, because signatures must be updated for the IDS to detect an attack

Behavior-based IDS is more complex than knowledge-based IDS and functions by attempting to “learn” normal user behavior patterns and then alarm when activity occurs that is outside of the normal use. Behavior-based IDS is sometimes referred to as anomaly-based IDS. There are benefits and drawbacks to a behavior-based IDS, with some of the benefits being

- ◆ Systems can dynamically respond to new, original, or unique exploits and attacks.
- ◆ Not dependent on specific operating systems.

Some drawbacks are

- ◆ High false alarm rates are very common. A recent evaluation in *NetworkWorld* blasted these systems for having such a high incidence of false alarms that real attacks are masked by the sheer volume of alarms.
- ◆ In environments where the usage patterns of the users and network resources are frequently changing, the IDS is unable to establish the baseline of “normal” behavior upon which to base any deviations.

IDSs have earned a bit of a mystique as being the silver bullet needed to prevent attacks before they become an issue. While the potential is certainly there, the reality is that the technology is not a substitute for a human being actively monitoring and managing network resources. Instead, an IDS is simply another tool in the well-prepared security professional’s toolbox.

Intrusion Response

Intrusion response is the principle of defining how to respond when an intrusion is detected. Intrusion response is often defined as part of the responsibilities of a Computer Incident Response Team (CIRT).

The primary responsibility of the CIRT is to define and execute the company's response to an incident via a process known as Incident Response Management. The CIRT response consists of the following:

- ◆ Coordinate how the notification and distribution of incidents should occur. There should be a defined escalation path to avoid situations of “forward this to everyone in your email list” from occurring.
- ◆ Mitigate the risk of an incident by minimizing disruptions and the costs involved in remediating the incident.
- ◆ Assemble teams of people to investigate and resolve potential incidents.
- ◆ Provide active input in the design and development of the company security policy.
- ◆ Manage and monitor logs.
- ◆ Manage the resolution of incidents, including post mortems of incidents.

Network Address Translation

Network Address Translation (NAT) is quite possibly the most misunderstood technology when it comes to security. This is largely due to a handful of vendors pushing it as a security mechanism, when it was not designed that way and does not function in that manner. Let's look at what NAT does and does not do.

NAT was designed, originally, to address the issue of IP address depletion on the Internet by allowing addresses on one network (referred to as the inside network) to be translated to a new address on a different network (referred to as the outside network). Typically, a company would have a range of addresses that were officially registered Internet IP addresses that they would use for communication with other Internet devices, but would run on different IP addresses (typically using the CIDR “non-routable” or “private” address ranges of 10.0.0.0, 172.16.0.0–172.31.0.0 or 192.168.0.0) on the internal network. In a classic NAT scenario, each internal address is translated to a unique external address. A variation of NAT is PAT (Port Address Translation) which instead of performing a one-to-one mapping of addresses,

performs a one-to-many mapping, using unique port numbers to differentiate between hosts. This allows a company to provide external access to hundreds of internal hosts while only using a single external IP address. Although this works great for allowing internal hosts to gain access to external resources, how does NAT work with providing external hosts access to internal resources? Let's look at that.

NAT can also be used to provide access to internal resources when used in conjunction with *policy routing*. This is sometimes referred to as *inbound NAT*. When the administrator has identified an internal resource that needs to be accessed, she can create a NAT table entry that maps the externally used IP address (the IP address external users are going to attempt to connect to) to the internally used IP address (the IP address of the system providing the service). As previously mentioned, PAT can also be used with inbound NAT to map a single external address to multiple internal addresses, using port numbers to differentiate between hosts. For example, a user connecting to IP address 1.1.1.1 on port 80 (HTTP) might be translated to the Web servers internal IP address, whereas a request to IP address 1.1.1.1 on port 21 (FTP) would be translated to a completely different server.

Because NAT can effectively hide the IP addresses that are being used internally, and because those internal addresses are often on the private range of addresses, NAT can provide a degree of security (and I use that term very lightly) to a network. More than security though, what NAT really provides is a boundary between networks. Unfortunately, because of the ability of NAT to mask the internal addresses, many folks believe that NAT alone is enough to secure a network from external risks. Nothing could be further from the truth. Section 9 of RFC 2993 sums up the problem of security and NAT very well:

NAT (particularly NAPT) actually has the potential to lower overall security because it creates the illusion of a security barrier, but does so without the managed intent of a firewall. Appropriate security mechanisms are implemented in the end host, without reliance on assumptions about routing hacks, firewall filters, or missing NAT translations, which may change over time to enable a service to a neighboring host. In general, defined security barriers assume that any threats are external, leading to practices that make internal breaches much easier.

—<http://www.ietf.org/rfc/rfc2993.txt?number=2993>

Although NAT can provide a measure of defense, NAT does not protect against things such as spoofed addresses and malicious content. As a result, NAT should never be considered a security solution, but rather it should be considered a component of a security solution, to be used in conjunction with firewalls and proxies.

Another issue to be aware of with NAT is the incompatibility of many types of encryption. NAT functions by transparently receiving packets destined for a host, and then building a new packet that it sends on behalf of the original request. When the response comes back, the device translates the response back and sends it to the original requestor. However, with many encryption methods, manipulation of the data is not permitted, and thus when NAT builds the new packet to send, the destination rejects it because the encryption is not correct anymore. One way around this problem is to configure the device doing NAT to not perform conversion functions on the packets. This works particularly with PPTP-based encryption. Another alternative, and gaining in popularity, is for the devices performing the encryption to encapsulate the encrypted data in TCP or UDP before sending it. As a result, NAT is performed on the TCP or UDP data, leaving the originally encrypted data unchanged. There are currently a number of RFCs under consideration to provide a standard method of dealing with the issues related to NAT and encryption (particularly IPSec).

Transparency

Transparency is simply the ability of a device to not appear to exist. Transparency can be a very effective security mechanism for a simple reason: How can you exploit something that does not appear to exist? In normal communications, when a device receives a packet for a service that is not running, the device notifies the sender that the service is not available. When a device is configured for transparency, though, rather than responding, “service is not running,” the device silently drops the packet, often forcing the sender to wait for a time-out period before it can attempt to connect again. Because the sender did not receive a response one way or the other, the sender is unable to determine what, if anything, might exist on the given IP address. Transparency is often used on firewalls to prevent connections to the external interface other than for services and addresses specifically advertised.

Another method of transparency is to configure a device to receive packets, but to not be able to send. This is typically used for IDS (Intrusion Detection System) deployment, so that the IDS can capture any potentially dangerous traffic without being susceptible to attack. One way of doing this is to cut the transmit pins of the network interface, thereby physically preventing the ability of the host to transmit. This is more of a sledgehammer method of transparency, however. A more subtle method is to configure an interface as a “probe” interface, but to not configure the interface with an IP address. Because the interface is going to receive any signals, whether it has an IP address or not, it can still receive data for processing against the IDS signature database. The device is incapable of processing any data for itself because it effectively has no protocol stack to use.

Hash Totals

Hashing is the process of assigning a value to represent some original data string. The value is known as the hash total. Hashing provides an efficient method of checking the validity of data by removing the need for the systems to compare the actual data, but instead allowing them to compare the value of the hash, known as the hash total, to determine if the data is the same or different. The hash value is represented in a database of some form, which allows for quicker indexing and searching for the original value. If the hash totals match, the data is the same. If the hash totals differ, the data is different. One of the best examples of this is Windows authentication. A common misconception is that when a user attempts to log on to a system, his username and password are sent to a domain controller for validation. Actually, the client generates a hash total based on the password the user enters, and sends that to the domain controller. The domain controller then checks the hash total against the hash total it has to represent the password. If they match, the user is allowed to log in. If they do not match, the user is prompted that he cannot log in. Because the actual password is never transmitted, there is an additional degree of security imparted in the transmission.

Email Security

As recent news articles would suggest, the importance of email security is becoming more and more of an issue. Not only is the securing and reliability of the mail datastore important, but the security of the content during transmission is equally important.

Security of email during transmission is pretty much the exclusive domain of encryption. Even if the email is able to be captured, the content is secure unless the content can be decrypted.

Another aspect of email security is securing the servers responsible for handling email. One of the biggest problems on the Internet today is the occurrence of UCE (unsolicited commercial email), better known as spam. Spam is bulk mail sent to people throughout the world. Most spam is not sent from the spammer's system, however. That would imply that the spammer would need to pay money for the bandwidth and servers that the spam requires for transmission. Instead, spammers look for SMTP servers that permit *relaying* of mail. Relaying is the capability of the SMTP server to send mail on behalf of someone else, in this case the spammer. To prevent your systems from being able to be used in this manner, ensure that you turn off relaying on the server. Now you might ask why you should do this, if it doesn't affect you. Well, aside from consuming your bandwidth, servers that leave their relays open tend to get added to various "blacklists" of Internet servers. If configured, other email servers will not accept email from or allow email to blacklisted servers.

Any discussion of email would be remiss without a discussion of viruses. Email is the single biggest method of spreading viruses today. Unfortunately, the best defense against viruses is the hardest thing to accomplish—educating users. As a result, it is critical to employ virus detection and removal software to detect and clean potentially harmful software. Even this stops short of effective protection however, because most virus software can only detect and remove viruses that it knows about. To augment the use of virus protection software, it is also recommended to block certain high-risk attachment types from being sent via email. A small list of attachment types to block would be scripts, executables, and files that contain macros (for example, Microsoft Word documents).

Although this might seem to be a hassle, it is far less of a hassle than having to clean up a virus breakout.

Facsimile and Printer Security

One of the biggest benefits of networking is the ability to share printer resources among multiple users. Unfortunately, this can also be a drawback in terms of security, because these resources tend to be located in shared common areas. The same is true for facsimile transmissions. A common method of securing this information is by giving each user her own printer and fax machine; this, however, might be cost prohibitive. An alternative is to have a secure fax machine and printer in a locked room with restricted access. If someone prints or faxes something, a designated person retrieves the information. I once worked for a place that had a secured printing area, but every user had access to the room. That effectively defeated the purpose of a “secure” area. The key to the security is not the area itself, but restricting access, with only designated persons allowed access to the area. Of course, any discussion of facsimile or printer security would be remiss if we did not mention a secure method of disposing of old papers, namely shredding of paper waste. Because of the advent of software that can actually put shredded documents back together, similar in fashion to a jigsaw puzzle, burning paper waste is the only effective method to ensure that the information cannot be recovered.

Common Attacks and Countermeasures

There are six classifications of network abuse (though this is by no means an exhaustive listing). Each of these is discussed in the following sections.

Class A Abuses

Class A network abuse is the result of unauthorized network access through the circumvention of security access controls. This is sometimes referred to as *logon abuse* and can range from legitimate users trying to access resources that they are not allowed to, to external threats attempting to gain access to a network. There are a number of techniques and countermeasures for class A network abuse:

- ◆ **Social engineering**—Social engineering is one of the hardest attacks to defend against due to the fact that the only real defense is user education. Social engineering can be as subtle as someone calling the help desk claiming that she lost her password and needs it reset; or someone calling users, pretending he's in the IT department, and saying he needs the user's password to test something. Masquerading is defined as a user pretending to be another user. Masquerading is often used as a component of a social engineering attack. Education is the key to preventing social engineering attacks.
- ◆ **Brute force**—Brute-force attacks tend to revolve around password cracking and hacking attempts. In a brute-force attack, the attacker is simply devoting all her resources toward gaining access to the system through a trial-and-error process. In the case of password cracking, she might repeatedly attempt to log on to a system hoping to gain access to the system. The defenses against brute force attacks depend on the type of attack. In the case of password attacks, a good password policy requiring the use of at least three character types (for example, uppercase, lowercase, and numbers) can help deter a hacker's ability to quickly guess the password. Protecting systems against brute force attacks is best done by ensuring that systems are adequately and timely patched against known vulnerabilities.

Class B Abuses

Class B network abuse is defined by non-business use of systems. This can be as surreptitious as someone printing personal items on company resources to as bold as visiting unauthorized Web sites. The most effective way to counteract class B network abuse is by way of a defined acceptable use policy (AUP) and an enforceable security policy with consequences for non-business use of resources. Content filtering and application proxies can also be used to provide a single point at which restrictions against unauthorized access can be enforced. Here are the types of Class B abuses:

- ◆ **PBX fraud and abuse**—PBX fraud costs companies millions of dollars every year. I know of a case in which a company had two employees in different countries. These employees were dating and racking up long-distance charges of up to \$5,000 per month, calling each other on the company's phone system. Several things can be done to prevent PBX fraud and abuse.

First, implement security on the phone system so that only authorized personnel can make long-distance phone calls. Second, ensure that each user must enter a unique code to gain access to make long-distance phone calls. Third, audit the phone calls made by users (identified by the user code) to detect suspicious use. Another form of PBX fraud involves an external user calling the company and asking to be connected to a long-distance number. This is a somewhat popular scam that can be defended against by educating the user community to not fall for such tactics.

- ◆ **Email and Internet abuse**—This is another area that costs companies millions of dollars every year, especially in regards to virus propagation and defense. From visiting inappropriate Web sites, to sending and receiving non-business-related emails and attachments, email and Internet abuse can be a very problematic issue to deal with. Many times employees believe that the emails they send are private. The AUP should make it clear that this is not the case. In addition to the AUP, email content filtering and virus-scanning software should be used to protect company resources. In terms of Internet abuse, the best defense is the use of proxies and Internet monitoring and blocking software to ensure that employees are only able to access resources that the company deems appropriate.

Class C Abuses

Class C network abuse is identified by use of eavesdropping techniques. These techniques can be active or passive in nature and include everything from listening to what someone is saying to tapping into a network to intercept network traffic. Some techniques for eavesdropping are

- ◆ **Network sniffing**—Capturing passing packets. As mentioned previously, network sniffing can provide the watcher with all the information they could need to compromise a system. One of the ways to defend against network sniffing, although not a complete solution, is through the use of switches for a network infrastructure. The most effective countermeasure though is through the use of encryption—for example, IPSec—because data that cannot be decrypted cannot be read.

- ◆ **Dumpster diving**—A social engineering technique, Dumpster diving is simply going through the trash to see if you can find something of value. This has been proven in a U.S. court of law to be an acceptable practice. The most effective defenses against dumpster diving are shredding and burning of trash.
- ◆ **Keystroke recording**—Keystroke recording can be used to capture all data entered into a computer. Because the program must be executed to capture data, a host-based IDS or similar system that can identify permitted programs and executables can be run to prevent the keystroke capturing program from executing.

Class D Abuses

Class D network abuse is identified by denial of service saturation of network services and resources. There are many types of denial of service attacks, but here are a few of the more popular:

- ◆ **SYN flooding**—As part of TCP communications, the devices attempting to communicate must synchronize the manner in which they will communicate. In a SYN flood, the server is inundated with requests to open a session, but the session is never completed. The server must wait for the establishment timeout to occur to clear the partial session, during which time it continues to be inundated with requests for more sessions. Eventually, the server runs out of resources with which to manage sessions and stops responding. SYN floods can be defended against by employing an IDS to detect and respond to SYN attacks. Additionally, the timely application of patches (a common theme) can also help to prevent SYN floods from being successful. Finally, increasing connection queue size and decreasing establishment timeouts can also prevent SYN floods from being successful.
- ◆ **Buffer overflows**—Buffer overflows are generally the result of poorly written and tested code. Buffer overflows can be exploited by performing actions that cause the system to run out of resources with which to service legitimate requests or sending excessive data that the system is unable to process properly. In some of the worst cases, buffer overflows can actually provide the ability to run arbitrary code on the affected system.

Chapter 2 TELECOMMUNICATIONS AND NETWORK SECURITY 151

The countermeasure to buffer overflows, aside from better code review, testing, and vendor accountability, is to apply patches in a timely fashion.

- ◆ **Teardrop attacks**—Teardrop attacks refer to the use of overlapping IP fragments that can cause the affected system to reboot or halt. Teardrop attacks can be addressed by applying patches to the affected systems.
- ◆ **LAND attacks**—LAND attacks are based on sending a device a packet that has the same source and destination IP address of the device that is being attacked. Ensuring that your devices are patched against LAND attacks is the best way to protect against them.
- ◆ **SMURF attacks**—This always brings visions to my mind of little blue men wreaking havoc on a network. In reality, a SMURF attack uses ICMP to spoof ICMP echo requests to a network broadcast address. This causes all the systems to respond to the spoofed address, saturating it with requests. The best way to defend against SMURF attacks is to prevent IP directed broadcasts on your routers and to configure your operating systems not to respond to packets sent to an IP broadcast address.
- ◆ **Distributed denial-of-service (DDoS) attacks**—A relatively new method of attacking, a DDoS uses hundreds or even thousands of hosts to inundate a device with more requests than it can handle. Considered a brute-force method of attack, a DDoS simply saturates the network link or server with more data than it has bandwidth or resources to handle. Unfortunately, the only real defense against a DDoS is to patch the systems (known as zombies) that are used to perpetrate the attack in an effort to prevent them from being used to launch a DDoS in the first place. After all, if there are not systems that can be used to launch a DDoS, it is not possible to cause a DDoS. Once a site is under attack, the only effective countermeasure is for the upstream neighbor to filter the unwanted traffic off of the circuit, try to determine where the attack is coming from, and notify those administrators so that they can stop the systems from continuing to execute the attack. A well-implemented DDoS can be a very difficult problem to deal with, mostly because the most effective defense is the responsibility of someone else (the admin of the zombies).

Class E Abuses

Class E network abuse is generally defined by network intrusion and prevention. As with DoS attacks, there are many types of intrusion to be aware of:

- ◆ **Spoof attacks**—Spoof attacks, or spoofing, is simply the process of an attacker appearing to be something other than it is. The goal is to attempt to get traffic delivered to a host that the hacker has control of. One of the more common spoof attacks is an ARP redirect in a switched environment. As you may recall, ARP is used to determine the MAC-to-IP addresses associations to allow for network communications. Using an ARP redirect attack, the hacker configures a system to claim to have a MAC address belonging to another system (typically the default gateway). When the switch receives traffic destined for the default gateway, it actually forwards the frame to the host performing the ARP redirect, because that is where the switch thinks the MAC address is located. The hacker can then run a packet sniffer to capture the data, and forward the frame to the default gateway, ensuring that the user never detects a problem. One of the countermeasures against ARP redirects is to maintain static ARP mappings, or to use port-based security to only allow certain MAC addresses to be used on certain ports. Another option is to maintain a mapping of “important” MAC addresses, and monitor traffic to see if other devices claim to have that MAC address.
- ◆ **Trojans**—Trojans are software that an attacker installs on a system (for example, by emailing a “check out this great whack-a-mole game” message) that typically exists to provide remote control capabilities of the affected system. Trojans are typically disguised as some sort of useful program, which increases the odds of it being run. Some common Trojans are Subseven, NetBus and BackOrifice. Some countermeasures against Trojans are through the use of file integrity tools such as Tripwire that can detect when files are added or modified and notify the administrator. Additionally, many commercial virus detection programs include the ability to detect and clean Trojans. In the case of Trojans that provide remote control capabilities, a good countermeasure is the use of egress filtering on your routers and firewalls. Egress filtering is the process of restricting all outbound traffic, only allowing the specific outbound traffic that users require.

Chapter 2 TELECOMMUNICATIONS AND NETWORK SECURITY 153

Although many security professionals will spend great time and detail performing ingress filtering—keeping people out, they generally do not spend as much time looking at specifically what needs to go out. Egress filtering should be a standard security countermeasure employed on all networks.

- ◆ **Viruses and worms**—Viruses and worms are perhaps the most dangerous daily form of network abuse there is. While viruses and worms often have the same effect on systems, a key difference between them is a worm’s ability to replicate, particularly by self-replication, its way around a network. Viruses, on the other hand, rely on some other form of distribution method (for example, a user emailing it or saving it to a floppy drive). The best defense against viruses and worms is the use of antivirus software and the timely application of updated signatures and patches. The signatures should be updated frequently. On-access and periodic system scans should be defined as a component of the security policy. In addition, the means to “push” a virus update to systems on the network is critical, as a new virus or worm is not going to wait until Friday when the new virus signatures are installed to take effect.
- ◆ **Back doors**—Back doors are mechanisms that an attacker places on a system that he can use to regain access to a system in the event that it is lost. The only real countermeasure against back doors is the complete rebuilding—what I call FDISK, FORMAT, REINSTALL—of a system that has been compromised. Although this is often an extremely painful and time-consuming process, it is the only effective way to ensure that a compromised system is not longer at risk. Some will contend that if they know they were hacked by something, they can simply “undo what was done” to get the system back. My response to this is, “How do you know that was the only thing you were hacked by?” If a system was compromised once that you caught, it was compromised 100 times that you have not yet caught.
- ◆ **TCP hijacking**—TCP uses sequence numbers to determine the state of the communication stream. With some TCP implementations, particularly certain Microsoft implementations, the sequence numbers used are not randomly determined, which allows an attacker to insert traffic into the data stream and “hijack” the session. This can cause the attacked computer to start responding to the attacker’s system thinking that it is the original trusted system.

- ◆ **Piggy-backing**—Piggy-backing refers to the process of using a legitimate user's connection to gain access to a system. This could be the result of a user leaving a connection open or incorrectly logging off. A countermeasure to piggy-backing is to use security policies (for example, Microsoft Group Policy) to enforce desktop timeouts and the locking of unused desktops.

Class F Abuses

Class F network abuse refers to probing attacks. A variation of eavesdropping, probing attacks are used by malicious users to gain information about a network in preparation of a network intrusion or other attack. Depending on the information able to be gathered, probe attacks can give an intruder a list of services and resources available on the network, and can even provide a diagram of the network layout and how systems are interconnected. There are a number of types of probes:

- ◆ **Port scans**—Port scans are used to query a system to determine the ports, and thus applications running, that are responding on a system. Port scans can be used to provide a rather in-depth list of services in use. A countermeasure to port scans is to only run the required services on devices and to configure systems to ignore requests for services that it is not running, as opposed to responding that the service is not there. This can cause scans to take significantly longer because the scanner needs to wait for a timeout to occur, thus increasing the likelihood of catching the scanner in the act.
- ◆ **Banner abuse**—Many services use banners that include information about the type of system the service is running on. Examples are HTTP, FTP, and SMTP banners. This information can be used to determine the types of exploits to which a system might be vulnerable. For example, if the FTP banner of a server tells me that the server is running on Microsoft Windows, I know what types of vulnerabilities the system might be susceptible to. A countermeasure for banner abuse is changing the banner to reflect something other than what the system is running, or to use proxies for external access to resources, thus preventing the prober from being able to communicate with the target host directly.

- ◆ **Sniffing**—This topic has been discussed in depth previously, but it is worth mentioning again as it provides an excellent mechanism with which to gather information, including things such as routing tables, MAC addresses, and network server lists that can then be spoofed.

FAULT TOLERANCE AND DATA RESTORATION

Reliability of data storage and access is a critical need for many businesses today. Reliability of data storage can often be handled through the use of redundant array of inexpensive disks (RAID). RAID uses multiple hard drives and differing fault tolerant scenarios to ensure that data loss does not occur in the event of disk failure. There are five levels of RAID:

- ◆ **RAID 0**—Used to provide a performance increase by allowing simultaneous read and writes through striping of data across multiple disks, RAID 0 provides no fault tolerance. If one disk fails, the data on all disks is lost.
- ◆ **RAID 1**—Better known as *mirroring*, RAID 1 duplicates the data on one disk to another disk. RAID 1 is a fairly expensive solution due to the fact that it requires double the storage; because it is a one-to-one duplication, it has 50% overhead (50% of the disks are not used for other than backup).
- ◆ **RAID 2**—Uses multiple disks and parity information; however, it has been replaced by other technologies and is no longer used. Parity tracks whether data has been lost or overwritten by use of a parity bit. The parity bit is calculated by calculating a group of data and measuring the bits set to 1. If the number of bits set to 1 is even, the parity bit is set to 1. If the number of the bits is odd, it is set to 0. When the data is read, if any bit data has been lost, the parity bit can be read to see whether the data sum should have an odd or even result, and the parity bit can then be changed to effectively re-create the data.
- ◆ **RAID 3–4**—RAID 3 performs byte-level striping and RAID 4 performs block level striping across multiple drives. Parity information is stored on a specific parity drive.

- ◆ **RAID 5**—By far the most popular fault tolerance method, RAID 5 stripes data and parity across all drives using interleave parity for data re-creation. Because reads and writes can be performed concurrently, RAID 5 offers a performance increase over RAID 1.

What happens if the entire server fails though? This is where the use of clustering technologies comes into play. There are two types of clustering concepts:

- ◆ **Data clustering**—Data clustering is the classic redundant server scenario. In this scenario, the administrator configures two servers as mirrors of each other, both sharing access to a common storage system. In the event that one of the servers fails, the services running on that server can be transferred to the backup server, hopefully with little to no impact on the user.
- ◆ **Network services clustering**—Also known as *load balancing*, network services clustering is used to improve system performance by distributing network requests among multiple servers which typically have the same functionality. The classic scenario is Web services, where each server maintains an exact copy of the Web site, thus allowing any of the servers to service client requests. If one of the servers is busy servicing a client request, another one can service it, and if one of the servers fails, the other servers can handle requests.

Even if you do everything possible to ensure that your email (or other critical data) is as fault tolerant as possible, sometimes everything fails and you are left with a molten puddle of goo and RAM chips. When this occurs, it is critical to have a backup of the data lost so that it can be re-created if possible. A number of backup methodologies are in use today:

- ◆ **Full backup**—The easiest backup methodology to manage, a full backup saves every file, every time. While easy to manage (you simply go to the last tape to restore everything), a full backup requires significant overhead in terms of the time it takes to back up the data and the cost of maintaining so many tapes that it can become cost prohibitive.

- ◆ **Incremental backup**—Incremental backups are an effective method of mitigating the risk between full backups by only backing up the data that has been changed or added recently. This can cut down significantly on the time and space required to backup a system. Incremental backups are generally performed between weekly full backups. In order to restore, you need to restore the last full backup, and then restore any incremental tapes leading up to the time of failure.
- ◆ **Differential backup**—A differential backup backs up files that have changed since the last full backup. As files are changed, they are added to the list of files to backup. The benefit to this approach is that a restore only requires the last full backup and the last incremental backup; however, it may take more time and tapes to backup than an incremental, depending on the amount of data changing.

Several types of backup media can be used:

- ◆ **Digital audio tape (DAT)**—Compact in size and low in cost, DAT backups are a very common method of backing up data. Unfortunately, DAT drives can back up only about 40GB per tape.
- ◆ **Quarter-inch cartridge (QIC)**—QIC is a backup tape system that is relatively small; only about 50GB per tape with more common systems supporting about 8GB per tape, and a fairly slow backup rate.
- ◆ **8MM tape**—8mm tape is an older tape storage system that has been largely replaced by DLT.
- ◆ **Digital linear tape (DLT)**—DLT is a 4mm tape system that provides a large capacity, currently up to 320GB per tape, and is extremely fast, with some systems backing up as fast as 120GB per hour.
- ◆ **CD/DVD**—With the advent of CD-R and DVD-RAM, more and more people are using them for desktop and small server backups and restores.
- ◆ **Zip**—Developed by Iomega, Zip drives have long been used as a desktop backup system with capabilities of backing up to 250MB. Zip drives use a media similar to a floppy disk.

- ◆ **Tape array**—A tape array is a cluster of 32–63 tape drives employed in a RAID fashion to provide for increased throughput and capacity.
- ◆ **Hierarchical storage management (HSM)**—HSM is a policy management methodology for backing up and restoring data in an enterprise. It is based on the principle that older data does not need to be restored as frequently, and thus can be moved to slower backup media to make room for new data on the faster backup systems. This is a critical technology in large enterprises where the cost of storing and maintaining backups can be huge.

Managing Network Single Points of Failure

Single points of failure on your network can make all of the server side fault tolerance and reliability implemented pretty much worthless. In the following sections, we take a look at cabling and topology failures, and how they affect the network.

Cable Failures

Cable failures are one of the most common types of network failures. Each cabling type has different vulnerabilities and effects, as illustrated in the following:

- ◆ **Coax cable**—Coax cable creates a single point of failure in the event that the cable is broken in any fashion. If there is a cable break, all systems on that cable will be unable to communicate.
- ◆ **Twisted pair**—Twisted pair, particularly unshielded twisted pair, is highly susceptible to interference. Twisted pair also has a shorter distance limitation than other cable types. On the high side though, only the device connected will be affected by the cable failure.
- ◆ **Fiber-optic cable**—Fiber-optic cable is immune to the electromagnetic interference that both coax and twisted pair are susceptible. Fiber optic has much longer distance limitations and is very fault tolerant, provided it is protected. The biggest problem with fiber optic is damage to the glass core, but if properly installed this should virtually never occur.

Topology Failures

One of the beauties of many topology failures is that with a well-designed network, the failure can be addressed and worked around via the use of redundancy of design. Each type of network topology has different vulnerabilities and effects, as illustrated in the following:

- ◆ **Ethernet**—Ethernet is the most popular network topology in use, largely because it can be implemented to be very tolerant of network failures. This is especially true in star, wired, and partially meshed hybrid designs.
- ◆ **Token-Ring**—Token-Ring was designed to be more fault tolerant than even Ethernet when implemented properly. Unfortunately the cost of a well-designed Token-Ring topology can be the biggest hindrance to good fault tolerance.
- ◆ **Fiber Distributed Data Interface (FDDI)**—Similar in design to Token-Ring, FDDI uses redundant rings to ensure that if the primary ring fails, devices can continue to communicate via the secondary ring.
- ◆ **Leased lines**—Leased lines provide a point-to-point connection and can be a single point of failure because they generally have no fault tolerance built into them. Effectively, when a leased line fails you are at the mercy of the provider to fix it in a timely fashion. A method of getting around this issue is to use redundant leased lines that are provided by different providers. Some networks will even use technologies such as ISDN to provide on-demand connections in the event of a leased-line failure.
- ◆ **Frame relay**—Frame relay is one of the most fault tolerant topologies because it was designed so that if any segment of the public network fails, traffic is diverted to other network segments. Fault tolerance can be further augmented by using multiple providers, similar to how leased lines work.

CASE STUDY: CODERED

ESSENCE OF THE CASE

The following points are the essence of the case:

- ▶ Microsoft has a vulnerability in its Web server software.
- ▶ Three worms were written that exploited this vulnerability.
- ▶ The worms spread by using commonly permitted traffic types, SMTP, HTTP, and TFTP to locate and infect other systems.
- ▶ The worms would deface legitimate Web sites.
- ▶ The worms would launch a DoS attack against a certain IP address.
- ▶ The worms would further expose systems by opening administrative access on the systems using the guest account.
- ▶ Due to the nature of the replication pattern, the worms could act as a DoS against network infrastructure equipment, particularly routers.

SCENARIO

In the late summer and early fall of 2001 there was a series of worms that were released that wreaked havoc on Windows-based computer systems throughout the world. These worms were known as CodeRed, CodeRedII, and Nimda.

CodeRed was a fairly complex worm that was discovered on July 16, 2001. While many worms prior to CodeRed were spread through using email, CodeRed was actually spread using the HTTP protocol. The worm was written to exploit a known vulnerability in Microsoft Web server documented in Microsoft security bulletin MS01-033. The worm functioned by exploiting a buffer overflow in the file IDQ.DLL which is part of the Microsoft Index Server product. CodeRed then potentially did a number of things:

- It would attempt to spread itself by attempting to connect systems on randomly determined subnets.
- On U.S. English systems it would deface the Web pages, causing them to display

```
Welcome to http:// www.worm.com !  
Hacked By Chinese!
```
- It would attempt to launch a Denial of Service attack against the IP address 198.137.240.91, which was the `www.whitehouse.gov` server (which has since been changed), by sending junk data to port 80.

CodeRed also had the unwitting effect of executing a denial of service against many routers as a result of its attempts to spread itself. What would happen is that the requests to port 80 on remote subnets would have to be handled by routers. In many cases, however, the requests were going to IP addresses that did not exist.

CASE STUDY: CODERED

What would happen then is that the routers would issue an ARP request for the IP address, because the router had no way of knowing that the IP address did not exist on the network. All the ARP requests would cause the router to fill its buffers waiting for responses that were never coming, thus preventing real data from passing through the router. This DoS was particularly effective in poorly designed networks that used class B address spaces that were largely vacant. Part of the beauty of CodeRed is that it really hit the wild on a Friday, which meant that many sites were totally unprepared for it due to the weekend.

On August 4, 2001, a variant of CodeRed known as CodeRedII was released. CodeRedII used the exact same exploit, and amazingly there were a lot of companies that were affected by it as well. CodeRedIII was much more aggressive than CodeRed, however, and it devoted more resources to distributing itself, making the effect of the router DoS much more severe. While the payload of CodeRed was ultimately the DoS attack against the White House Web site, CodeRedII was designed to deploy a Trojan on infected systems that provided full remote control and execution capabilities. This Trojan effectively provided a back door for access to the Web server.

On September 18, 2001, a new worm hit the streets known as Nimda (which is admin spelled backwards). Nimda exploited a Web Folder Traversal exploit that allowed a hacker to create a URL that would provide access to any directory structure and files on the server. This exploit was documented in a Microsoft security bulletin released on October 17, 2000. One of the most shocking things about Nimda was the aggressiveness with which it attempted to spread itself.

Although we had seen worms use email or HTTP as a delivery mechanism, Nimda was one of the first to use both techniques at the same time, as well as using TFTP to spread, which was a fairly new method. Nimda did a number of things:

- Nimda used mass mailing to spread itself.
- Nimda modified numerous files, which allowed it to be run anytime any of the modified files were executed.
- Nimda would create a large amount of files, which could cause a system to run out of disk space.
- Nimda opened a significant security hole by sharing the c: drive. In conjunction, Nimda added the user "Guest" to the local administrators group, allowing for anyone to then connect to the share as an administrator.

Numerous variants of Nimda have since been released, but all are fundamentally the same in function.

ANALYSIS

So what can we learn from these three worms?

First and foremost in importance is an examination of the date of infection and the date the exploits were documented by Microsoft. In each case, Microsoft had released a patch before the worm had hit the street. In the case of Nimda, the patch was released almost a full year prior to the creation of Nimda. The lesson is to apply patches from vendors in a timely fashion.

Second, these worms were able to infiltrate companies because they allowed access into the network on ports that were generally insecure.

continues

CASE STUDY: CODERED*continued*

I actually had the privilege of using a tool that would play back traffic it recorded; it was running when CodeRed hit. At the company in question, we were able to play back and observe the CodeRed traffic entering the network through the VPN connections and through a couple of servers that were accessible via the Web and were able to access the internal network. In the case of Nimda, it required the ability to email an executable attachment in order to spread via email. Because many companies did not block executables from being emailed, the worm was able to spread very easily. The lesson? Implement good security perimeters and only allow the traffic that you absolutely need. Filter email traffic as well.

Third, buffer overflows are a bad thing. While the end user has little ability to directly address buffer flows, what the user does have is the ability to come down hard on vendors that do not test their code for overflow vulnerabilities before they release the software. Microsoft took a beating over these exploits and should have.

Fourth, run only those network services that are required. One reason that CodeRed was so successful is that virtually every version of Windows OS runs a Web server by default. Most administrators do not modify the installation; rather they just click Next, Next, Finish, and deploy the system regardless of whether it will ever actually be used as a Web server. Had more systems, particularly desktops and servers that did not host Web services, not had IIS installed, the impact of these worms would have been significantly less.

Ultimately though, the lesson lies not in blame for Microsoft over the fact that the exploits exist but in the realm of the security professional. What may well be the worst thing about CodeRed and Nimda is that patches existed well before the worms hit that protected against them and they were still able to spread like wildfire. As security professionals, we must be more vigilant.

CHAPTER SUMMARY

The Telecommunications and Network Security domain includes a massive amount of information to learn. It provides details on the processes, systems, and technologies that make up the backbone of networking and network security.

We started with an examination of the OSI model and the benefits of using layered architectural design models. We looked at the processes that occur at each layer, and the protocols that enable them.

Next, we examined the characteristics and topologies that define a network. We examined the cable types and interconnectivity models that define a network. We built upon that to examine the roles that the network technologies of Ethernet, Token-Ring, and ARCnet have, and looked at the network devices that tie everything together.

We looked at the role of firewalls on the network and the different types of firewalls that exist. We also looked at the architectures that firewalls are deployed in, with an examination of the pros and cons of each type of design.

After establishing network functionality and security of the LAN, we covered WAN connection methods and the technologies and devices that facilitate WAN communication. Next, we took a look at how to provide remote access to a network, with an examination of dial-up and VPN-based solutions. We examined how site-to-site and client-based VPNs work and how we can secure our VPNs using PPTP, L2TP, and IPSec.

Our look at telecommunications networks wrapped up with an exploration of TCP/IP and an examination of the four-layer DoD architectural model. We finished up with a look at the protocols that enable TCP/IP communications.

With networking out of the way, we proceeded into a discussion of the security needs of a network. We discussed the CIA triad and the goals of confidentiality, integrity, and availability in regards to our systems and services. We examined how to incorporate security boundaries into our network designs and the benefit of separating resources and users.

KEY TERMS

- 802.2
- 802.3
- 802.5
- Authentication
- Bridge
- Broadcast
- Denial of service
- Eavesdropping
- Encryption
- Extranet
- Firewall
- Gateway
- Hub
- IPSec
- L2TP
- Layer 3 switch
- Linear bus
- Mesh
- Multicast
- Multiplexor
- NAT
- OSI
- Packet analyzer

continues

CHAPTER SUMMARY *continued*

- PEM
- PPP
- PPTP
- Proxy
- RAID
- Ring
- Router
- SDLC
- SET
- SKIP
- SLIP
- S/MIME
- SSL
- Star
- SWIPE
- Switch
- TCP/IP
- TLS
- Token-Ring
- Tree
- Tunnel
- Unicast
- VLAN
- VPN

Next, we looked at the Trusted Network Interpretation concepts and how to use the TCSEC to ensure that your systems and processes are secure and functional.

A discussion of security protocols followed, where we looked at the network, transport, and Application layer security protocols that can be used to protect data. We also looked at network monitoring and packet sniffers and how they can be used for good and bad to determine what the data on the network is doing.

We also took a detailed look at the components that make up an Intrusion Detection System and compared network- and host-based IDS as well as knowledge-based and behavior-based IDS. We wrapped up the intrusion discussion with a look at intrusion response.

We looked at the functions and pitfalls of NAT as a security solution, and then took a look at forms of network abuse and how to defend against them. The chapter wrapped up with a look at fault tolerance and assuring data availability.

As mentioned, this domain has a wide reach and covers a lot of ground. However, if you apply the principles of a layered design to studying these concepts, separating them into easy-to-learn pieces, the totality of the information will be much easier to digest.

APPLY YOUR KNOWLEDGE

Exercises

2.1 Designing Network Topologies

You are the administrator of a new network. The CIO has tasked you with the responsibility of designing the corporate network while providing the maximum degree of security. The following requirements have been given:

- ◆ The internal network must be secured against external and internal threats.
- ◆ Several servers will need to be accessed by external users. The internal network must be secured, even if these servers are compromised.
- ◆ Traveling and home-office users will need access to internal network resources.
- ◆ Outbound Internet (WWW) access must be screened and filtered.

Estimated Time: 1 hour

1. Design and diagram a network topology that will meet these needs.
2. To design the most effective solution, let's review the requirements:

- The internal network must be secured against external and internal threats.

The most effective device to implement for securing a network against external threats is a firewall. The most effective method to secure against internal threats is to use an IDS (both network and host based) and to ensure that all systems are properly patched and running virus protection.

- A number of servers will need to be accessed by external users. The internal network must be secured, even if these servers are compromised.

The type of firewall and firewall design to implement differ with every network. Because multiple servers will need to be accessed by external resources, a screened subnet firewall design would be preferred. The screened subnet and internal firewall will protect the internal network in case the externally accessible servers are compromised. Likewise, the external firewall will provide some degree of security for the externally accessible servers against external threats. This can be handled via the use of a circuit proxy/stateful inspection firewall. This will provide an excellent combination of speed and security. Using an IDS to monitor traffic entering and exiting the screened subnet in both directions will further protect against security compromises.

- Traveling and home office users will need access to internal network resources.

The most effective method to provide internal access is through the use of VPN connections. It is also recommended to use multiple screened subnets, one for the externally accessible servers and one for the VPN connections. This allows you to manage each group of external traffic separately, as well as providing a single point to block VPN access if required.

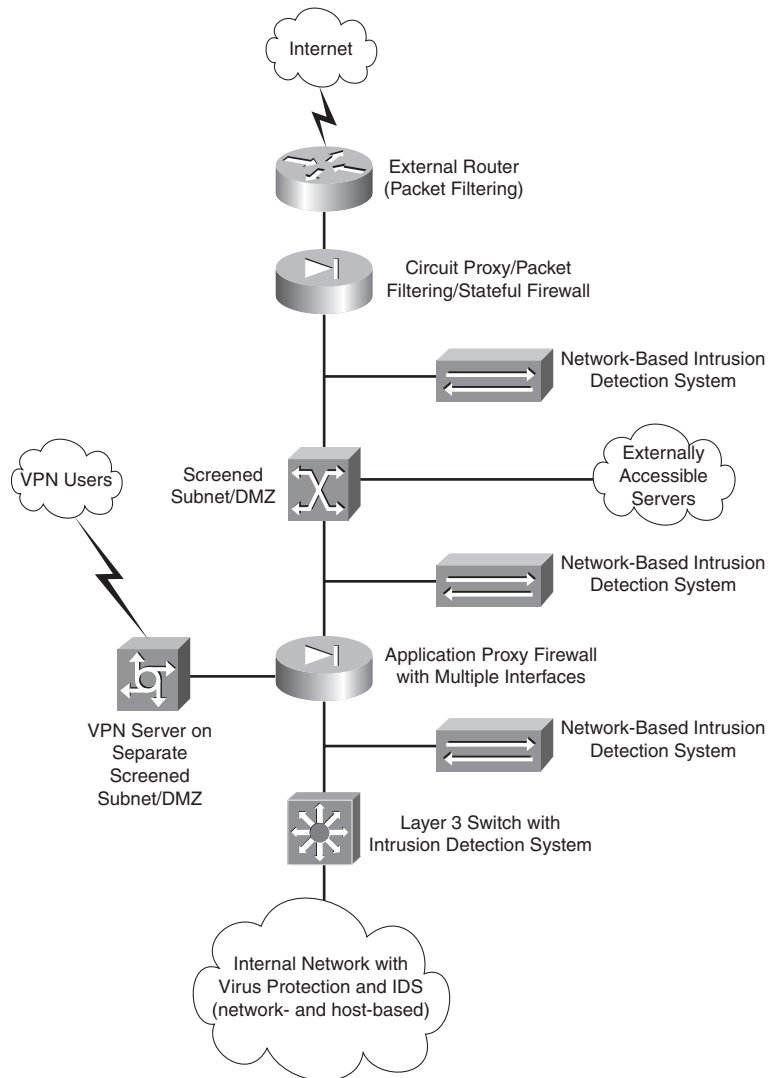
- Outbound Internet (WWW) access must be screened and filtered.

The most effective method to screen outbound Internet access is through the use of an application proxy. This can be provided through the use of application proxy firewalls. This is an effective internal firewall choice in our screened subnet design.

Figure 2.20 illustrates how the solution comes together.

APPLY YOUR KNOWLEDGE

Diagram of exercise 1 solution.



Review Questions

1. What are some of the benefits of a layered architecture model?
2. What are the six firewall types and what are their characteristics?
3. How can a network administrator provide secure remote connections to the network?
4. What are the differences between authentication and encryption?

APPLY YOUR KNOWLEDGE

5. What are the six classifications of network abuse, and what are their characteristics?
6. How can a network administrator increase the reliability of network data?

Exam Questions

1. Which OSI layer is primarily responsible for negotiating dialog control between systems and applications?
 - A. Application layer
 - B. Transport layer
 - C. Session layer
 - D. Internet layer
2. Routers are devices which function at which layer of the OSI Model?
 - A. Data Link layer
 - B. Internet layer
 - C. Physical layer
 - D. Network layer
3. Coaxial cable is typically used in which LAN topology?
 - A. Mesh
 - B. Linear bus
 - C. Star
 - D. Tree
4. What is the minimum UTP cable specification that supports transmitting of data at 100Mbps speeds?
 - A. Category 3
 - B. Category 5
 - C. Category 5e
 - D. 10BASE-T
5. What is the single point of failure in a star topology?
 - A. The cable
 - B. The computer
 - C. The hub or switch
 - D. The NIC
6. Which device is responsible for separating broadcast domains?
 - A. Router
 - B. Switch
 - C. Bridge
 - D. Repeater
7. What is used at the Data Link layer for the delivery of data to hosts?
 - A. IP address
 - B. IPX address
 - C. ARP
 - D. Hardware address

APPLY YOUR KNOWLEDGE

8. Ethernet uses which access method?
 - A. Carrier Sense, Multiple Access/Collision Avoidance
 - B. Token passing
 - C. Carrier Sense, Multiple Access/Collision Detection
 - D. LAN emulation
9. Sending and receiving data at the same time is an example of which type of communication?
 - A. Simplex
 - B. Multicast
 - C. Full-Duplex
 - D. Half-Duplex
10. A device that keeps track of the connection state of conversations is known as a(n) _____?
 - A. Application proxy
 - B. NAT device
 - C. Stateful inspection firewall
 - D. Packet filtering firewall
11. Using a perimeter network to secure internal resources from external sources, while still providing limited access to devices on the perimeter network is an example of a _____?
 - A. Packet filtering firewall design
 - B. Screened subnet firewall design
 - C. Screened host firewall design
 - D. Dual homed host firewall design
12. T1 lines are typically used for which type of WAN connection?
 - A. Circuit-switched
 - B. Cell-switched
 - C. Remote access
 - D. Dedicated
13. CHAP and PAP Authentication can be used with which type of technology?
 - A. HDLC
 - B. X.25
 - C. Dedicated WAN connections
 - D. PPP
14. What is used as the underlying connection for establishing a VPN connection?
 - A. Dial-up remote access
 - B. The Internet
 - C. Circuit-switched connections
 - D. Dedicated connections
15. What is used for providing connection-oriented delivery in the TCP/IP protocol suite?
 - A. SNMP
 - B. UDP
 - C. IP
 - D. TCP
16. What does ARP do?
 - A. Resolves known IP addresses to MAC addresses
 - B. Resolves known MAC addresses to IP addresses

APPLY YOUR KNOWLEDGE

- C. Resolves NetBIOS names
D. Resolves hostnames
17. What TSCEC division specifies that discretionary protection through the use of auditing occurs?
- A. Division A
B. Division B
C. Division C
D. Division D
18. SWIPE provides security at which layer?
- A. Physical
B. Transport
C. Application
D. Network
19. S/MIME is used to secure which type of data?
- A. Web traffic
B. IPX
C. Email
D. Database queries
20. A device that examines network traffic to look for anomalies from the normal traffic patterns is an example of a(n) _____?
- A. Application proxy firewall
B. Stateful packet inspection firewall
C. Behavior-based IDS
D. Host-based IDS
21. Social engineering is an example of what class of network abuse?
- A. Class A
B. Class B
C. Class C
D. Class D
22. Class D network abuse is identified by what?
- A. Non-business use of systems
B. Denial of service
C. Network intrusion
D. Probing

Answers to Review Questions

1. There are three primary benefits to using a layered reference model:
 - It divides the complex network operation into smaller, easier-to-manage pieces or layers.
 - It facilitates the ability to make changes to the functions and processes at one layer without needing to make changes at all layers.
 - It defines a standard interface for multi-vendor integration. By using a standard interface, the details of how a particular layer functions are hidden from all the other layers.

See “The OSI Layers” section for more information.

APPLY YOUR KNOWLEDGE

2. The following are the six types of firewalls and a brief description of their characteristics:
 - **Packet filtering**—Packet-filtering firewalls are similar in use and function to routers. In fact, many routers include packet-filtering capabilities. Packet-filtering firewalls function by comparing received traffic against a rules set that defines what traffic is permitted and what traffic is denied.
 - **Application proxy**—Application-filtering firewalls function by reading the entire packet up to the Application layer before making a filtering decision. Whereas a packet-filtering firewall generally cannot differentiate between the valid application data and invalid application data, the application proxy firewall can.
 - **Circuit proxy**—Circuit proxy firewalls are a bit of a hybrid between application proxies and packet-filtering firewalls. With a circuit proxy, the firewall creates a circuit between the source and destination without actually reading and processing the application data. In that sense, it is a proxy between the source and destination. However, because it does not actually process the application data, it is functionally like a packet filter.
 - **Stateful inspection**—All firewalls being considered today should perform *stateful packet inspection*. When a host sends a packet to the destination, the destination is going to process the data and potentially send a response. This network connection state is tracked by the firewall and then used in determining what traffic should be allowed to pass back through the firewall. Because these firewalls can examine the state of the conversation, they can even monitor and track protocols that are otherwise considered connectionless, such as UDP or certain types of remote procedure call traffic.
 - **Dynamic packet filtering**—A dynamic packet filtering firewall is generally used for providing limited support of connectionless protocols such as UDP. It functions by queuing all the UDP packets that have crossed the network perimeter, and based on that will allow responses to pass back through the firewall.
 - **Kernel proxy**—Kernel proxy firewalls are typically highly customized and specialized firewalls that are designed to function in kernel mode of the operating system. This provides for modular, kernel-based, multi-layer session evaluation using customized TCP/IP stacks and kernel-level proxies.
3. See the “Firewalls” section for more information.
3. Secure remote connections and access to the network can be provided through the use of VPN connections. A good VPN connection will use both authentication and encryption to ensure that only permitted connections are allowed to be established and that all the data transmitted is encrypted for security. See the “Providing Remote Access” and “VPNs (Virtual Private Networks)” sections for more information.
4. Authentication is a process in which the identity of the remote host is validated. Encryption is a process in which the data transmitted is secured so that it can be read only by the correct destination host. A secure network data delivery process combines both authentication and encryption to validate the source and destination systems and protect the integrity of the data. See the “Wireless” section and “Network Layer Security” section for more information.

APPLY YOUR KNOWLEDGE

5. The following are the six classifications of network abuse, and a brief description of their characteristics:
- **Class A abuses**—Class A network abuse is the result of unauthorized network access through the circumvention of security access controls. This is sometimes referred to as logon abuse, and can range from legitimate users trying to access resources that they are not allowed to, to external threats attempting to gain access to a network.
 - **Class B abuses**—Class B network abuse is defined by non-business use of systems. This can be as surreptitious as someone printing personal items on company resources to as bold as visiting unauthorized Web sites.
 - **Class C abuses**—Class C network abuse is identified by the use of eavesdropping techniques. These techniques can be active or passive in nature and include everything from listening to what someone is saying to tapping into a network to intercept network traffic.
 - **Class D abuses**—Class D network abuse is identified by denial of service saturation of network services and resources.
 - **Class E abuses**—Class E network abuse is generally defined by network intrusion and prevention.
 - **Class F abuses**—Class F network abuse refers to probing attacks. A variation of eavesdropping, probing attacks are used by malicious users to gain information about a network in preparation of a network intrusion or other attack.

Depending on the information that can be gathered, probe attacks can give an intruder a list of services and resources available on the network, and can even provide a diagram of the network layout and how systems are interconnected.

See the “Common Attacks and Countermeasures” section for more information.

6. Reliability of network data can be best assured through the use of fault-tolerant systems and data recovery methods. Some examples of fault-tolerant systems are the use of RAID to protect data-storage systems and clustering to provide fail-over redundancy. If the network administrator is unable to prevent the failure, the use of data backup and recovery systems can further provide data reliability. See “Fault Tolerance and Data Restoration” for more information.

Answers to Exam Questions

1. **B.** The Session layer is responsible for negotiating dialog control between systems and applications. The Application layer is responsible for interfacing to the user. The Internet layer is not an OSI layer. The Transport layer is responsible for end-to-end communications. See “Session Layer” for more information.
2. **D.** Routers function at the Network layer of the OSI model. Switches and bridges function at the Data Link layer. The Internet layer is not an OSI layer, but routers could be considered Internet layer devices. Hubs and repeaters function at the Physical layer. See “Network Layer” for more information.

APPLY YOUR KNOWLEDGE

3. **B.** Coaxial cable is typically used in a linear bus topology. Mesh, star, and tree topologies are typically created with UTP cabling. See “Coax” and “Linear Bus Topology” for more information.
4. **B.** Category 5 is the minimum UTP specification that will run at 100Mbps. Category 3 is not capable of transmitting data at 100Mbps. Although category 5e is capable of transmitting at 100Mbps, it is not the minimum specification. 10BASE-T is not a UTP cable specification. See “Unshielded Twisted Pair” for more information.
5. **C.** The hub or switch is the single point of failure in a star topology. Cable failures in a star topology affect only the devices connected to that cable. Computer or NIC failures affect only the device in question. See “Star Topology” for more information.
6. **A.** Routers are responsible for separating broadcast domains. Switches and bridges will forward broadcasts, potentially creating broadcast storms in a looped network. Repeaters repeat every signal, regardless of what it is. See “Network Layer” and “Routers” for more information.
7. **D.** The hardware address is used at the Data Link layer for delivering data to hosts. IP and IPX addresses are used for logical addressing at the Network layer. ARP is used to resolve IP addresses to MAC addresses; it is not used for the delivery of data. See “Data Link Layer” for more information.
8. **C.** Ethernet uses Carrier Sense, Multiple Access/Collision Detection for its access method. Token passing is used in FDDI and Token-Ring networks. Carrier Sense, Multiple Access/Collision Avoidance is used for Arcnet. LAN Emulation is used for ATM networks. See “Ethernet” for more information.
9. **C.** Full-duplex allows a system to send and receive data at the same time. Simplex is unidirectional transmission only. Multicast is an addressing method that allows multiple hosts to receive the same data. Half-duplex is a bidirectional transmission method, however it can only transmit in one direction at a time. See “Ethernet” for more information.
10. **C.** A stateful packet inspection firewall keeps track of the connection state of conversations. Application proxies process the data packet to verify that it is the proper application data. NAT devices simply translate addresses. Packet filtering firewalls do not track connection state; they simply forward or filter based on access lists. See “Firewalls” for more information.
11. **B.** A screened subnet firewall design protects internal resources by using a perimeter network, while providing external access to devices on the perimeter network. A packet-filtering firewall design does not contain a screened subnet. In a screened host firewall design the exposed host is on the internal network, not on a perimeter network. A dual-homed host firewall uses a host that is connected to the external and internal network; however, it will not forward packets between those networks. See “Firewalls” for more information.
12. **D.** T1 lines are typically used for dedicated WAN connections. Circuit-switched connections are typically used for dial-up and backup connections. Cell-switched connections are used in ATM. Remote access is not a WAN access connection as much as it is an access method. See “Dedicated Connections” for more information.

APPLY YOUR KNOWLEDGE

13. **D.** CHAP and PAP authentication is used with PPP. HDLC and X.25 do not use authentication. Dedicated WAN connections is not a valid response. See “Point-to-Point Protocol and Serial Line Internet Protocol” for more information.
14. **B.** The Internet is used as the underlying connection for establishing VPNs. Dial-up remote access is when the user dials into the corporate network directly. Circuit-switched and dedicated connections are WAN connection methods, not VPN connection methods. See “Virtual Private Networks” for more information.
15. **D.** TCP is used for providing connection-oriented communications in the TCP/IP protocol suite. SNMP is used for managing IP devices. UDP and IP are connectionless. See “Transport Layer Protocols” for more information.
16. **A.** ARP resolves a known IP address to an unknown MAC address. RARP resolves known MAC addresses to unknown IP addresses. WINS resolves NETBIOS names. DNS resolves host names. See “Internet Layer Protocols” for more information.
17. **C.** TSCEC division C specifies that discretionary protection through the use of auditing should occur. Division A uses formal security verification to ensure security. Division B specifies that mandatory access rules exist. Division D uses minimal protection, if any. See “Trusted Network Interpretation” for more information.
18. **D.** SWIPE provides Network layer security. Application layer security is provided through protocols such as S/MIME and PEM. SSL and TLS are protocols that provide Transport layer security. Physical layer security can be provided by controlling access to the physical cabling. See “Network Layer Security Protocols” for more information.
19. **C.** S/MIME is used to provide security for email data. Web traffic is secured via HTTPS and SSL. IPX can be secured by encapsulating it in other protocols. Database queries can be secured by application/Presentation layer encryption or encapsulating it in other protocols such as IPsec. See “Application Layer Security Protocols” for more information.
20. **C.** A behavior-based IDS looks for anomalies in traffic patterns. An application proxy firewall proxies connections between hosts and examines the application data to ensure integrity. Stateful packet inspection firewalls track conversation state to determine whether to permit or deny traffic. Host-based IDSs run on and monitor an individual host. While a host-based IDS might be a behavior-based IDS, it does not have to be one. See “Intrusion Detection” for more information.
21. **A.** Social engineering is an example of Class A network abuse. Class B network abuse is indicated by abuse of network resources. Class C network abuse is indicated by the use of eavesdropping. Class D network abuse is indicated by a denial of service or saturation of network resources. See “Common Attacks and Countermeasures” for more information.
22. **B.** Class D network abuse is identified by denial of service. Non-business use of systems is an example of Class B network abuse. Network intrusion is an example of Class E network abuse. Probing is an example of Class F network abuse. See “Common Attacks and Countermeasures” for more information.

APPLY YOUR KNOWLEDGE**Suggested Readings and Resources**

1. Comer, Douglas. "Internetworking with TCP/IP," Volume 1. In *Principles, Protocols, and Architecture*, Prentice Hall, 2000.
2. Stevens, W. Richard. *TCP/IP Illustrated, Volume 1*. Addison Wesley, 1994.
3. <http://standards.ieee.org/802news/802july2002.html> (newsletter on IE802 Working Groups).
4. <http://standards.ieee.org/getieee802/portfolio.html?agree=ACCEPT> (802 standards documentation available for free online).
5. <http://www.iana.org/assignments/port-numbers> (a list of all registered TCP and UDP port numbers).
6. <http://www.icsalabs.com/html/communities/firewalls/index.shtml> (firewall testing criteria, FAQs, and whitepapers that can provide more detailed information about firewalls).
7. <http://www.itsecurity.com/tutor/detectingcabletaps.htm>
8. <http://www.rfc-editor.org/rfcsearch.html> (search the RFC index or find and read RFCs by number or subject).
9. RFCs:
 - IP: 791
 - TCP: 3168
 - UDP: 768
 - ICMP: 792
 - TLS: 2246
 - PEM 1421, 1422, 1423, 1424
 - NAT: 2993

OBJECTIVES

Understand the principles of security management.

- ▶ In understanding information security management, there are a number of principles you need to know to create a managed security program. These principles go beyond firewalls, encryptions, and access control. They are concerned with the various aspects of managing the organization's information assets in areas such as privacy, confidentiality, integrity, accountability, and the basics of the mechanisms used in their management.

Know what management's responsibility is in the information security environment.

- ▶ Management cannot just decree that the systems and networks will be secure. They must take an active role in setting and supporting the information security environment. Without management support, the users will not take information security seriously.

Understand risk management and how to use risk analysis to make information security management decisions.

- ▶ Managing security is the management of risk. Knowing how to assess and manage risk is key to an information security management program.



CHAPTER

3

Security Management and Practices

OBJECTIVES

Know how to set policies and how to derive standards, guidelines, and implement procedures to meet policy goals.

- ▶ Policies are the blueprints of the information security program. From policies, you can set the standards and guidelines that will be used throughout your organization to maintain your security posture. Then, using those standards, you can create procedures that can implement the policies.

Set information security roles and responsibilities throughout your organization.

- ▶ From management to the users, everyone who has access to your organization's systems and networks is responsible for their role in maintaining security as set by the policies. Understanding these roles and responsibilities is key to creating and implementing security policies and procedures.

Understand how the various protection mechanisms are used in information security management.

- ▶ Protection mechanisms are the basis of the data architecture decision that will be made in your information security program. These are the basis for the way data is protected and provide a means for access.

Understand the considerations and criteria for classifying data.

- ▶ Protecting data is the objective of every information security program. Therefore, we look at how that data can be classified so it can be securely handled.

Determine how employment policies and practices are used to enhance information security in your organization.

- ▶ Even with the press concentrating on the effects of denial-of-service attacks and viruses, the biggest threats come from within. Improving on the employment policies and practices to perform better background checks and better handle hiring and termination, as well as other concerns to help minimize the internal threat, are important information security practices.

Use change control to maintain security.

- ▶ One of the jobs of a Trojan horse is to replace a program with one that can be used to attack the system. Change control is one defense against this type of attack. Using change control to maintain the configuration of programs, systems, and networks, you can prevent changes from being used to attack your systems.

Know what is required for Security Awareness Training.

- ▶ The best security policies and procedures are ineffectual if users do not understand their roles and responsibilities in the security environment. Training is the only way for users to understand their responsibilities.

OUTLINE

Introduction	179	Identify from Whom It Is Being Protected	209
		Setting Standards	209
		Creating Baselines	210
		Guidelines	210
		Setting and Implementing Procedures	210
Defining Security Principles	180	Examining Roles and Responsibility	212
CIA: Information Security's Fundamental Principles	180	Management Responsibility	213
Confidentiality	181	User Information Security Responsibilities	213
Integrity	182	IT Roles and Responsibilities	214
Availability	183	Other Roles and Responsibilities	214
Privacy	183	Understanding Protection Mechanisms	215
Identification and Authentication	184	Layering	216
Passwords	185	Abstraction	217
Nonrepudiation	188	Data Hiding	217
Accountability and Auditing	188	Encryption	217
Keystroke Monitoring	189	Classifying Data	218
Protecting Audit Data	190	Commercial Classification	219
Documentation	190	Government Classification	220
Security Management Planning	191	Criteria	221
		Creating Procedures for Classifying Data	221
Risk Management and Analysis	192	Employment Policies and Practices	222
Risk Analysis	194	Background Checks and Security Clearances	222
Identifying Threats and Vulnerabilities	195	Employment Agreements, Hiring, and Termination	223
Asset Valuation	196	The Acceptable Usage Policy	224
Qualitative Risk Analysis	202	Termination	224
Countermeasure Selection and Evaluation	203		
Tying It Together	204		
Policies, Standards, Guidelines, and Procedures	205		
Information Security Policies	206		
How Policies Should Be Developed	206		
Define What Policies Need to Be Written	207		
Identify What Is to Be Protected	207		

OUTLINE

Job Descriptions 225

Job Rotation 225

Managing Change Control 226

Hardware Change Control 226

Software Change Control 227

Security Awareness Training 227**Chapter Summary 228****Apply Your Knowledge 230****STUDY STRATEGIES**

- ▶ Even if you are not part of your organization's management team, watch how management works in the information security environment. Take the practices and strategies written here and look at not only how your organization implements them, but how they can be improved. This type of lateral thinking will help on the exam and can make you a valuable contributor to your organization's security posture.
- ▶ The notes throughout the chapter point out key definitions and concepts that could appear on the exam. They are also key components that all managers should understand.

This chapter covers Domain 3, Security Management Practices, 1 of 10 domains of the Common Body of Knowledge (CBK) covered in the Certified Information Systems Security Professional Examination. This domain is divided into several objectives for study.

“Security management entails the identification of an organization’s information assessment and the development, documentation, and implementation of policies, standards, procedures, and guidelines that ensure confidentiality, integrity, and availability. Management tools such as data classification, risk assessment, and risk analysis are used to identify the threats, classify assets, and to rate their vulnerabilities so that effective security controls can be implemented.

Risk management is the identification, measurement, control, and minimization of loss associated with uncertain events or risks. It includes overall security review, risk analysis, selection and evaluation of safeguards, cost benefit analysis, management decision, safeguard implementation, and effectiveness review.

The candidate will be expected to understand the planning, organization, and roles of the individual in identifying and securing an organization’s information assets; the development and use of policies stating management’s views and position on particular topics and the use of guidelines, standard, and procedures to support the policies; security awareness training to make employees aware of the importance of information security, its significance, and the specific security-related requirements relative to their position; the importance of confidentiality, proprietary, and private information; employment agreements; employee hiring and termination practices; and risk management practices and tools to identify, rate, and reduce the risk to specific resources.”

—Common Body of Knowledge study guide

INTRODUCTION

Security management can be difficult for most information security professionals to understand. It is the bridge between understanding what is to be protected and why those protections are necessary.

Using basic principles and a risk analysis as building blocks, policies can be created to implement a successful information security program.

As part of creating that program, information security management should also understand how standards and guidelines also play a part in creating procedures. When doing this, every user's role and responsibilities should be accounted for by understanding how to protect the organization's information assets.

The role of data as a significant part of the organization's information assets cannot be minimized. Data provides the fuel that drives your organization, but it is the asset that is the most vulnerable. Protecting this asset means understanding the various classifying mechanisms and how they can be used to protect your critical assets.

This chapter covers all these issues and discusses security awareness and managing people in your information security environment.

DEFINING SECURITY PRINCIPLES

To understand how to manage an information security program, you must understand the basic principles. These principles are the building blocks, or primitives, to being able to determine why information assets need protection.




FIGURE 3.1 Security's fundamental principles are confidentiality, integrity, and availability.

CIA: Information Security's Fundamental Principles

Remembering that information is the most important of your organization's assets (second to human lives, of course), the first principles ask *what is being protected*, *why*, and *how do we control access*? The fundamental goal of your information security program is to answer these questions by determining the *confidentiality* of the information, how can you maintain the data's *integrity*, and in what manner its *availability* is governed. These three principles make up the CIA triad (see Figure 3.1).

The CIA triad comprises all the principles on which every security program is based. Depending on the nature of the information assets, some of the principles might have varying degrees of importance in your environment.

Confidentiality

Confidentiality determines the secrecy of the information asset.

Determining confidentiality is not a matter of determining whether information is secret or not. When considering confidentiality, managers determine the level of access in terms of how and where the data can be accessed. For information to be useful to the organization, it can be classified by a degree of confidentiality.

To prevent attackers from gaining access to critical data, a user who might be allowed access to confidential data might not be allowed to access the service from an external access port. The level of confidentiality determines the level of availability that is controlled through various access control mechanisms.

Protections offered to confidential data are only as good as the security program itself. To maintain confidentiality, the security program must consider the consequences of an attacker monitoring the network to read the data. Although tools are available that can prevent the attacker from reading the data in this manner, safeguards should be in place at the points of transmission, such as by using encryption or physically safeguarding the network.

Another attack to confidentiality is the use of social engineering to access the data or obtain access. Social engineering is difficult to defend because it requires a comprehensive and proactive security awareness program. Users should be educated about the problems and punishments that result when they intentionally or accidentally disclose information. This can include safeguarding usernames and passwords from being used by an attacker.

Cryptography is the study of how to scramble, or encrypt, information to prevent everyone but the intended recipient from being able to read it. Encryption implements cryptography by using mathematical formulas to scramble and unscramble the data. These formulas use an external piece of private data called a *key* to lock and unlock the data.

Cryptography can trace its roots back 4,000 years to ancient Egypt where funeral announcements were written using modified hieroglyphics to add to their mystery. Today, cryptography is used to keep data secret. For more information on cryptography, see Chapter 5, “Cryptography.”

Integrity

With data being the primary information asset, *integrity* provides the assurance that the data is accurate and reliable. Without integrity, the cost of collecting and maintaining the data cannot be justified. Therefore, policies and procedures should support ensuring that data can be trusted.

Mechanisms put in place to ensure the integrity of information should prevent attacks on the storage of that data (*contamination*) and on its transmission (*interference*). Data that is altered on the network between the storage and the user’s workstation can be as untrustworthy as the attacker altering or deleting the data on the storage media. Protecting data involves both storage and network mechanisms.

Attackers can use many methods to contaminate data. Viruses are the most frequently reported in the media. However, an internal user, such as a programmer, can install a back door into the system or a logic bomb that can be used attack the data. After an attack is launched, it might be difficult to stop and thus affect the integrity of the data. Some of the protections that can be used to prevent these attacks are intrusion detection, encryption, and strict access controls.

Not all integrity attacks are malicious. Users can inadvertently store inaccurate or invalid data by incorrect data entry, an incorrect decision made in running programs, or not following procedures. They can also affect integrity through system configuration errors at their workstations or even by using the wrong programs to access the data. To prevent this, users should be taught about data integrity during their information security awareness training. Additionally, programs should be configured to test the integrity of the data before storing it in the system. In network environments, data can be encrypted to prevent its alteration.

Availability

Availability is the ability of the users to access an information asset. Information is of no use if it cannot be accessed. Systems should have sufficient capacity to satisfy user requests for access, and network architects should consider capacity as part of availability. Policies can be written to enforce this by specifying that procedures be created to prevent denial-of-service (DoS) attacks.

More than just attackers can affect system and network availability. The environment, weather, fire, electrical problems, and other factors can prevent systems and networks from functioning. To prevent these problems, your organization's physical security policies should specify various controls and procedures to help maintain availability.

Yet access does not mean that data has to be available immediately. Availability of information should recognize that not all data has to be available upon request. Some data can be stored on media that might require user or operator intervention to access. For example, if your organization collects gigabytes of data daily, you might not have the resources to store it all online. This data can be stored on an offline storage unit, such as a CD jukebox, that does not offer immediate access.

Privacy

Privacy relates to all elements of the CIA triad. It considers which information can be shared with others (confidentiality), how that information can be accessed safely (integrity), and how it can be accessed (availability).

As an entity, privacy is probably the most watched and regulated area of information security. Laws, such as the U.S. Federal Privacy Act of 1974, provide statutes that limit the government's use of citizens' personal data. More recently, the Health Insurance Portability and Accountability Act (HIPAA) authorizes the Department of Health and Human Services to set the security and privacy standards to cover processing, storing, and transmitting individual's health information to prevent inadvertent or unauthorized use or disclosure.

Laws and regulations have been difficult to keep up-to-date as the technology moves forward. The federal government has been able to keep up by using directives and mandates within the executive branch. However, this has not helped private industry. Regulations, such as those mandated by the U.S. Federal Trade Commission (FTC), attempt to help, but the FTC lacks enforcement capabilities.

If not mandated by law or regulation, organizations should look at the privacy of their own information assets. Aside from having to be concerned about the privacy of employee information, an organization needs to be concerned about the disclosure of customer information that might not be regulated.

Information collected through contact, such as via the Internet, does not require a privacy statement, but the FTC does say organizations should have one. That privacy statement should reflect how the data is handled and available to the users whose information is being collected.

Monitoring privacy has other concerns. Preventing the unauthorized disclosure of data might require monitoring of data transmission between systems and users. One area of concern is the monitoring of email. Email monitoring can include content monitoring to watch for unauthorized disclosure of information. However, before doing so, an organization must ensure that policies are in place that state what might be monitored or disclosed.

Finally, security professionals introduce an additional problem to the privacy of information because of their nearly unlimited access to all resources. Although we would like to think that all professionals have integrity, some have other agendas or lack the knowledge to prevent accidental disclosure. Security professionals should be limited to the information that is necessary to perform their tasks. Policies can be created to have additional checks and balances to ensure integrity of the data.

Identification and Authentication

Information security is the process of managing the access to resources. To allow a user, a program, or any other entity to gain access to the organization's information resources, you must identify them and verify that the entity is who they claim to be.

The most common way to do this is through the process of *identification and authentication*.

The process of identification and authentication is usually a two-step process, although it can involve more than two steps. Identification provides the resource with some type of identifier of who is trying to gain access. Identifiers can be any public or private information that is tied directly to the entity. To identify users, the common practice is to assign the user a username. Typically, organizations use the user's name or employee identification number as a system identifier. There is no magic formula for assigning usernames—it is a matter of your preference and what is considered the best way of tracking users when information appears in log files.

The second part of the process is to authenticate the claimed identity. The following are the three general types of authentication:

- ◆ What the entities know, such as a personal identification number (PIN) or password
- ◆ What the entities have, such as an access card, a smart card, or a token generator
- ◆ Who or what the entity is, which is usually identified through biometrics

Out of these general types of authentication, if two or more are used, the authentication is called *strong authentication*. For physical security, a user with an access card commonly must enter a PIN. For authentication to a system or network, a common method is to use a PIN or pass code with a token generator. Although biometrics is a way to identify who the entity is, another step is still necessary to strengthen the authentication.

Passwords

Of these methods, passwords and PINs are the most common forms of authentication. Although passwords become the most important part of the process, they also represent the weakest link. As a security manager, you must manage the process in such a way to minimize the weakness in the process.

NOTE

Understand the Principle of

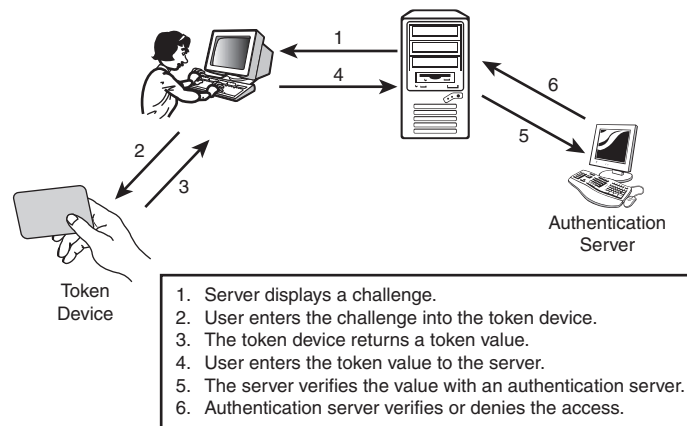
Authentication Authentication is a matter of what the entity knows, what they might have, or who the entity is. For strong authentication, use at least two of these principles.

Users typically create passwords that are easily guessed. Common words or the names of spouses and children leave the password open to dictionary or social engineering attacks. To prevent these attacks, some organizations use a password generator to create passwords that cannot be cracked using typical attacks. The problem is that these passwords are usually not that memorable, which causes the users to write them down, leaving them open to another type of social engineering attack in which another user finds the documented password.

Password management involves trying to create a balance between creating passwords that cannot be guessed and passwords users don't need to write down. Policies can mandate several strategies that can be effective in mitigating some of these problems. Following are some of the methods management should use when mitigating these problems:

- ◆ **Password generators**—These are usually third-party products that can be used to create passwords out of random characters. Some products can be used to create memorable passwords using permutations of random or chosen words or phrases.
- ◆ **Password checkers**—These are tools that check the passwords for their probability of being guessed. They are designed to perform typical dictionary attacks, and they use information on the system in an attempt to guess the password using social engineering. These checkers also use common permutations of these attacks, anticipating what a user might try. For example, users commonly use 0s in the place of the letter *o*. The strength of the password is determined by how many attempts the tool makes to guess the password.
- ◆ **Limiting login attempts**—These can prevent attackers from trying to log in to systems or prevent networks from using exhaustive attacks. By setting a threshold for login failures, the user account can be locked. Some systems can lock accounts for a period of time, whereas others require administrator intervention.
- ◆ **Challenge-Response**—These are also called *cognitive passwords*. They use random questions that the user would provide the answer to in advance or use a shared secret. When the user logs in, the system picks a random question that must be answered successfully to gain access. This is commonly used on voice response systems (for example, social security number, account number, ZIP code, and so on) and requires the answer to more than one challenge.

◆ **Token devices**—These are a form of one-time password authentication that satisfies the “what you have” scenario. Token devices come in two forms: synchronous and asynchronous. A *synchronous* token is time-based and generates a value that is used in authentication. The token value is valid for a set period of time before it changes and is based on a secret key held by both the token (usually a sealed device) and the server providing authentication services. An *asynchronous* token uses a challenge-response mechanism to determine whether the user is valid. After the user enters the identification value, the authentication server sends a challenge value. The user then enters that value into the token device, which then returns a value called a *token*. The user sends that value back to the server, which validates it to the username. Figure 3.2 demonstrates these steps.



NOTE

PKI Using public key or asynchronous encryption technologies requires the use of a public key infrastructure (PKI) to manage the process.

FIGURE 3.2

Authentication using an asynchronous token device.

◆ **Cryptographic keys**—These combine the concepts of “something you have” and “something you know.” Using public key cryptography, the user has a private key (or digital signature) that is used to sign a common hash value that is sent to the authentication server. The server can then use the known public key for the user to decrypt the hash. To strengthen the authentication process, the user is asked to enter a PIN or passphrase that is also added to the hash to strengthen the authentication process.

NOTE

Understanding Nonrepudiation

Nonrepudiation is the ability to ensure the authenticity of a message by verifying it using the message's digital signature. Remember, digital signatures require a certificate to generate the signature and a PKI to save the public key for when the message is verified.

Nonrepudiation

Nonrepudiation is the ability to ensure that the originator of a communication or message is the true sender by guaranteeing authenticity of his digital signature. Digital signatures are used not only to ensure that a message has been electronically signed by the person who purported to sign the document, but also to ensure that a person cannot later deny that he furnished the signature.

One way to authenticate the digital signature is to verify it with the public key obtained from a trusted certification authority (CA). When used in PKI, the CA stores the public key that could be used to verify the signature. However, digital signatures might not always guarantee nonrepudiation. One concern is the trust of the signature and the CA. For example, some commercial CA products do not require verification of the person buying the signature but trusts that his credit card is valid. In pretty good privacy, you have to trust the signers of the user's certificate.

Regardless of how your organization tries to implement nonrepudiation, there will be some risk based on the trust of the information used for validation. Biometric verification can help in the process, but that means you must trust the certification process.

Accountability and Auditing

With the user authenticated to the system and network, most administrators use the various audit capabilities to track all system events. Systems and security administrators can use the audit records to

- ◆ Produce usage reports
- ◆ Detect intrusions or attacks
- ◆ Keep a record of system activity for performance tuning
- ◆ Create evidence for disciplinary actions or law enforcement

Accountability is created by logging the events with the information from the authenticated user, which might also include date, time, network address, and other information that could further identify the condition that caused the event. Events are audited through system and network facilities designed to help monitor from the lowest levels. These facilities also have Application Program Interfaces (APIs) that can allow applications to audit pertinent event information.

Administrators can set up auditing to capture systems events. However, if you set up auditing to capture everything, you will create logs that can take up all available disk space. Rather, you should set a parameter defining a *threshold*, or *clipping level*, of the event to be logged. Setting thresholds is typical in the configuration of intrusion detection systems (IDSs). An IDS has the tendency to log a lot of erroneous events called *false positives*. Setting thresholds can cut down on the number of errors logged.

The auditing of systems requires active monitoring and passive protections. *Active* monitoring requires administrators to watch the ongoing activities of the users. One way this can be done is via keystroke monitoring. *Passive* monitoring is done through the examining of audit data maintained by each system. Because the audit data is usually stored on the system, it should be protected from alteration and unauthorized access. These auditing principles are discussed in the following sections.

Keystroke Monitoring

Keystroke monitoring is a type of audit that monitors what a user types. It watches how the user types individual words, commands, or other common tasks and creates a profile of that user's characteristics. The keystroke monitor can then detect whether someone other than the profiled user tries to use the system.

Another form of keystroke monitoring is the capture of what the user types. These types of keystroke monitors capture some of the basic user input events, allowing forensic analysis of what the user is doing. This is a more controversial form of auditing because it has been used by law enforcement in recent high-profile cases.

In either case, there are two problems with this type of auditing:

- ◆ The generation of a lot of data
- ◆ Privacy issues

Because of the nature of the data captured, no clipping level can be set. Therefore, you must ensure that there is enough storage for all the captured information to be stored.

Privacy issues are a concern in all types of monitoring, but especially with keyboard monitoring. Unless used by law enforcement with the proper authorization, you should ensure that your organization has the proper policies in place and users have been notified of those policies.

NOTE

Magic Lantern The FBI has been looking at new ways of doing covert investigation of criminals on the Internet. One tool they use is called Magic Lantern. As a follow-up to the Carnivore program, the FBI covertly installs Magic Lantern on a targeted computer system to trap keystroke and mouse information. Magic Lantern has been used to break the encryption of a suspected criminal. As this is written, that case has yet to come to trial, but the constitutionality of the FBI using Magic Lantern will be a central question.

Otherwise, you run the risk of being accused of violating a user's civil rights and liberties. Although this has not been resolved in the courts, you should not try this without the proper policies in place because you do not know what would happen if the monitored user tried to test this in court.

Protecting Audit Data

There will come a time when your organization has to handle an incident. This incident can come from within your organization's network or from the Internet. The only way you will have to figure out how the incident occurred is through log analysis. However, the analysis of the logs can be only as successful as the integrity of the data.

Operating systems have many ways of maintaining the log data integrity, including the capability to store it across a network. Maintaining the integrity of the data is important for analysis. If the incident involves an attack, law enforcement can use the data gathered by the audits to investigate and prosecute the attacker. For the audit data to be used in legal proceedings, it must be proven that the integrity of the audit data has been maintained and there was no possibility for it to be altered. In the legal world, that is called *proving the chain of custody*. If the prosecutor cannot prove the chain of custody, the audit data cannot be used as evidence.

There are more reasons than law enforcement, but I put the emphasis on it because, if your protection procedures can pass that test, they will pass the others. It becomes important in any situation where legal proceedings might be involved, such as firing an employee for violating policies. Audit data used in the decision can be subpoenaed if the employee sues your organization, which requires the same chain of custody rules.

Documentation

When I talk to organizations about the condition of their security documentation, most admit that it is not up-to-date. Others say that it is too accessible because it details the controls and settings of various devices. In either case, documentation can become a weak link in the security chain. By not keeping up with documentation, there could be no explanation of how the controls are configured to satisfy policies, which would make their replacement in an emergency situation difficult.

Making the documentation accessible can be a controversial issue. Some believe that the more open security is, the better it can be reviewed and hardened. Review is one thing, but some people could use this information for unscrupulous purposes. If the user who has access to the full description of the security controls is also a disgruntled employee or even someone engaging in industrial espionage, it might be in your organization's best interest to restrict access to security documentation.

SECURITY MANAGEMENT PLANNING

Understand the principles of security management.

Planning for information security includes preparation to create information security policies that will be the guidance for the entire information security program. To create the policy, management should plan to perform a risk analysis on the information assets to be protected. The risk analysis will identify the assets, determine risks to them, and assign a value to their potential loss. Using this, management can make decisions on the policies that best protect those assets by minimizing or mitigating the risks.

The final aspect of information security management is education. Management is responsible for supporting the policy not only with its backing, but also by including policies and the backing for educating users on those policies. Through security awareness training, users should know and understand their roles under the policies. This is discussed further in the "Security Awareness Training" section, later in this chapter.

Managing an information security program changes with the release of every new operating system and with every new communications enhancement. Over the years, network technology has changed how information assets are protected. In the past, data was stored and accessed through mainframes where all the controls were centralized. Networked systems change this paradigm by distributing data across the network.

It does not help that network protocols were invented to share information and not with security in mind. In the beginning, security was left up to each system's manager in a small society of network users.

NOTE

Network's Importance to Security Management Network management is also important to security management. You should understand the roles of networks and some of the tools, such as virtual private networks (VPNs) and extranets.

As technology grew, the information assets became less centralized and management had the problem of maintaining the integrity of the network and the information being used on the systems on the networks. Although there is a move to try to centralize management of servers and information security, information security management needs to take into account everywhere the information assets touch.

Network computing has brought new paradigms to the sharing of information. Using technologies such as virtual private networks (VPNs) and extranets, organizations can forge new types of relationships based on sharing information assets. These partnerships have organizations connecting their networks to share information in a way that was unheard of as recently as 10 years ago. Managers planning these partnerships also should keep in mind how to maintain the security of other information assets not involved in those agreements. Both organizations should consider undergoing a risk analysis specific to the connectivity required for this partnership to provide appropriate protections.

RISK MANAGEMENT AND ANALYSIS

Understand risk management and how to use risk analysis to make information security management decisions.

Risk management is the process of assessing risk and applying mechanisms to reduce, mitigate, or manage risks to the information assets. Risk management is not about creating a totally secure environment. Its purpose is to identify where risks exist, the probability that the risks could occur, the damage that could be caused, and the costs of securing the environment. Even if there is a risk to information assets, risk management can determine that it would cost more to secure the asset than if it was damaged or disclosed.

Risk management is not as straightforward as finding the risk and quantifying the cost of loss. Because risks can come from varying sources, an information asset can have several risks. For example, sales data stored on a network disk has the risk of

- ◆ Unauthorized access from internal or external users
- ◆ Loss from a software or hardware failure
- ◆ Inaccessibility because of a network failure

Risk management looks at the various possibilities of loss, determines what would cause the greatest loss, and applies controls appropriately. As the risk manager, you might want to reduce all the risk to zero. This is a natural emotional reaction to trying to solve risk. However, you might find that it is impossible to prevent unauthorized access from internal users while trying to ensure accessibility of the data. Here, you must look at the likelihood of the risk and either look for other mitigations or accept it as a potential loss to the organization.

Assessing risk for information security involves considering the types of loss (*risk category*) and how that loss might occur (*risk factor*).

Risk Category

- ◆ **Damage**—Results in physical loss of an asset or the inability to access the asset, such as cutting a network cable.
- ◆ **Disclosure**—Disclosing critical information regardless of where or how it was disclosed.
- ◆ **Losses**—These might be permanent or temporary, including the altering of data or the inability to access data.

Risk Factor

- ◆ **Physical damage**—Can result from natural disasters or other factors, such as power loss or vandalism.
- ◆ **Malfunctions**—The failure of systems, networks, or peripherals.
- ◆ **Attacks**—Purposeful acts whether from the inside or outside. Misuse of data, such as unauthorized disclosure, is an attack on that information asset.
- ◆ **Human errors**—Usually considered accidental incidents, whereas attacks are purposeful incidents.
- ◆ **Application errors**—Failures of the application, including the operating system. These are usually accidental errors, whereas exploits of buffer overflows or viruses are considered attacks.

Every analyzed information asset has at least one risk category associated with one risk factor. Not every asset has more than one risk category or more than one risk factor. The real work of the risk analysis is to properly identify these issues.

NOTE

Risk Analysis Identifies a risk, quantifies the impact, and assesses a cost for mitigating the risk.

Risk Analysis

Risk analysis is a process that is used to identify risk and quantify the possible damages that can occur to the information assets to determine the most cost-effective way to mitigate the risks. A risk analysis also assesses the possibility that the risk will occur in order to weigh the cost of mitigation. As information security professionals, we would like to create a secure, risk-free environment. However, it might not be possible to do so without a significant cost. As a security manager, you will have to weigh the costs versus the potential costs of loss.

IN THE FIELD

BUSINESS VERSUS GOVERNMENT RISK ANALYSIS

A risk analysis for a government agency is no different from one performed for a nongovernment organization. The difference is how the information is used. Nongovernment entities can use the costs of mitigating the risk and the expected gain to determine whether to add countermeasures and which ones would be the most cost-effective. Most nongovernment entities work like this, including non-profit corporations.

Because of laws, regulations, and legislative oversight, government agencies (particularly on the federal levels) have to run in a risk adverse environment rather than a risk-managed environment. Thus, agencies provide security controls that minimize the risk to a zero-cost, regardless of the costs, to prevent them from being campaign fodder. It is why the government will spend more money to secure systems than a private corporation will.

On completion of the risk analysis, the information allows the risk manager to perform a cost-benefit analysis (CBA), comparing safeguards or the costs of not adding the safeguards. Costs are usually given as an annualized cost and can be weighed against the likelihood of occurrence. As a general rule, safeguards are not employed when the costs of the countermeasure outweighs the potential loss. For example, an information asset is worth \$10,000 should it be lost. Table 3.1 shows a possible analysis of this asset.

TABLE 3.1**BASIC RISK ANALYSIS ON A \$10,000 ASSET**

<i>Cost of Countermeasure</i>	<i>Gain/(Loss)</i>	<i>Analysis</i>
\$0	(\$10,000)	By doing nothing, if the asset is lost, there could be a complete loss that costs \$10,000.
\$5,000	\$5,000	If the countermeasure costs \$5,000, you will gain \$5,000 in providing the protection by mitigating the loss.
\$10,000	\$0	The cost of the countermeasure equals the cost of the asset. Here, you might weigh the potential for the countermeasure to be needed before making a decision.
\$15,000	(\$5,000)	With the countermeasure costing more than the asset, the benefit does not make sense in this case in terms of financial cost.

For information security planning, the risk analysis allows management to look at the requirements and balance them with business objectives and the costs. For an information security program to be successful, the merging of security processes and procedures with the business requirements is essential. A major part of that is the protection of the assets, and the risk assessment helps in that analysis.

Identifying Threats and Vulnerabilities

The previous section identified the various risk categories and factors that go into a risk analysis. For that analysis to weigh the potential for a risk to occur, the analysis should identify the threats and vulnerabilities that could occur.

There is no single way to identify whether a threat or vulnerability could occur in the environment being analyzed. Most environments are so complex that a vulnerability in one area could affect another area of the business. These *cascading errors* could be caused not only by a malicious attack, but also by errors in processing, which are called *illogical processing*.

NOTE

Threat Agents These are what cause the threats by exploiting vulnerabilities.

NOTE

Loss Potential This is what would be lost if the threat agent is successful in exploiting a vulnerability.

NOTE

Delayed Loss This is the amount of loss that can occur over time.

Identifying the threats to information assets is the process of identifying the *threat agents* that can cause a threat to the environment. Threat agents can be human, programmatic (such as an error or malware), or a natural disaster. The risk factors in the previous section provide a view into the number of possible threat agents an asset could have. Audits look at all the potential threat agents and determine which factors result in the risk to the asset.

After the threat agents, vulnerability, and risk have been identified, the risk analysis then concentrates on the *loss potential*, or what would be lost if the threat agent exploited the vulnerability. Whether the loss is from corruption or deletion of data to the physical destruction of computer and network equipment, there will be a cost to the loss of the asset. The loss is not limited to the cost of the asset. Risk analysis should also consider the loss of productivity, whether it be a delay or halt in work.

Not every loss will occur immediately. Take disclosure of critical data, for example. The loss from when the data is disclosed might not happen immediately. But if the disclosure was to a competitor involved in industrial espionage, the potential loss could occur over time in the form of lost clients and business. The loss potential for this type of delayed loss can attempt to estimate the costs to recover. Because the nature of the losses are unknown, making this type of estimate can be difficult.

Another delayed loss can be embedded in the cost of business. If data that is used to calculate fees, taxes, or other fiscal obligations is corrupted, a loss potential exists for interest and penalties that would have to be paid when the problems are discovered, which will be more than the costs to repair the damage. In more extreme cases, your organization could lose the confidence of its customers and investors, which could cause additional damage.

Asset Valuation

There are two ways to evaluate assets and the risk associated with their loss. The *quantitative* approach attempts to assign a dollar value to the risk for analyzing the cost of the potential effectiveness of the countermeasure. A *qualitative* approach uses a scoring system to rank threats and effectiveness of the countermeasures relative to the system and environment. Most commercial organizations prefer the quantitative approach because it allows for a way to plan budgets and for nontechnical management to understand the impact of their decisions.

NOTE

Quantitative Versus Qualitative A *quantitative* approach to risk analysis uses monetary values to assess risk. The *qualitative* approach uses a scoring system to determine risk relative to the environment.

However, a qualitative analysis is good for understanding the severity of the risk analysis relative to the environment, which is easier for some to understand.

When using the quantitative approach, you should remember that it cannot quantify every asset and every threat. When looking at the values at the extremes, whether high or low, the numbers tend to not reflect the reality of the quantitative analysis. It is up to the team doing the risk analysis to determine which approach is best.

IN THE FIELD

AN INTERNAL RISK ANALYSIS VERSUS USING OUTSIDE CONSULTANTS

Some might feel that their own systems and security professionals could perform the risk assessment. They do know the systems and understand the processing that occurs. However, although the people your company employs might be very competent, they might be too intimate with operations to be able to tell a technical risk from a process risk. Outsiders do not have the same ties, so they are not prejudiced by “what has been.”

When selecting an outside company to do a risk assessment, make sure it has the resources to understand the latest security information and industry best practices so it can provide a complete risk assessment. It must understand all the risks involved in all aspects of information technology. Because these companies do this on a daily basis, they have more insights into what to expect as they perform their tests.

Risk analysis is an investigation into the various assets, assigning risk and determining mitigations. To do this, the risk assessment team must investigate all the assets, taking into account all the variables that can affect the costs. The steps that are followed in a risk analysis are

1. Identify the assets.
2. Assign value to the assets.
3. Identify the risks and threats corresponding to each asset.
4. Estimate the potential loss from that risk or threat.
5. Estimate the possible frequency of the threat occurring.

6. Calculate the cost of the risk.
7. Recommend countermeasures or other remedial activities.

Each step is explained in Step By Step 3.1.

STEP BY STEP

3.1 Risk Analysis Steps

1. Identify the assets. When you identify your information assets, you must consider more than the systems and network components. Information assets can also be the organization's data. A company's sales data that contains customer information and buying habits is as much of an asset as the disk and systems that store the information. Risk analysts will look at the organization's business process and ask which information is important to the business processes. In this process, more emphasis can be put on the information that is important, such as sales data, rather than the company phone book.

This is where maintaining documentation and having a solid configuration management system can help. Rather than forcing a full discovery of all assets, including programs and databases, the documentation and configuration management systems can point to the bulk of the assets and provide a basis to begin the analysis. This is not to say that a risk assessment cannot be performed without this help. Some risk assessments are performed to gather this information, which is perfectly reasonable when establishing a new or more stringent information security program.

2. Next, you must assign value to the assets. Assigning value is not a simple task. For hardware or software, the value can be the purchase or the replacement costs. Setting the value to information assets is where the process becomes difficult. To determine value, you would answer the following questions:
 - How much revenue does this data generate?
 - How much does it cost to maintain?

Chapter 3 SECURITY MANAGEMENT AND PRACTICES 199

- How much would it cost if the data were lost?
- How much would it cost to recover or re-create?
- How much would it be worth to the competition?

3. After all the assets are identified, the analysis then identifies all the threats and risks. The various risk categories are examined, and the various factors are applied until a list of possible threats is created. There is no scientific way to determine which risk categories apply to an asset—it is a subjective determination. However, some common sense should prevail. For example, data cannot be damaged by fire, but the disks on which it resides can be. The risk for the data could be damage or unavailability because of hardware failure, which reduces a number of risk factors and potential countermeasures.

4. The next step is to go through the various assets and the threats to estimate how much would be lost if the threat occurs. Obviously, this is easy for hardware and software because costs can be taken from invoices or actual replacement costs. But what happens when the asset is data? How much would it cost if access to critical data were lost? How much would it cost to be recovered or regenerated? What if it was improperly disclosed?

When estimating the costs for the loss, all factors should be considered. For example, if workstations are infected with a virus, the cost of recovery should be counted, and so should the loss of productivity. Estimating productivity loss is not easy because the salaries and benefits for each employee affected should be considered, as well as the duration of the loss. Although a number of employees at different salary levels might work on the recovery effort, many times an estimate is based on an average salary. The numbers produced are appropriate for a risk analysis.

The estimated cost of the potential loss is used to calculate the single-loss expectancy (SLE) for the asset. SLE uses the asset value and the exposure factor (see step 5) to give the dollar amount of the potential loss if the threat came to pass. These calculations are discussed in step 6.

NOTE

Single-Loss Expectancy (SLE) This is the amount of the potential loss for a specific threat.

continues

continued

5. The frequency of occurrence is used to estimate the percentage of loss on a particular asset because of a threat. Also called the *exposure factor (EF)*, this value recognizes that a threat does not result in a total loss. For example, a fiber-optic cable running between two buildings being cut by a maintenance worker affects only the cable and the productivity for its cut, which might be only 20% of the organization's infrastructure. For this asset, the EF would be 0.20 for calculations.

Risk analysis is based on the loss over the course of a year. The *annualized rate of occurrence (ARO)* is the ratio of the estimated possibility that the threat will take place in a 1-year time frame. The ARO can be expressed as 0.0 if the threat will never occur, through 1.0 if the threat will always occur. For example, the ARO for a workstation virus might be set to 1.0, whereas a power outage to the network operations center that might occur once every 4 years would have an ARO of 0.25.

NOTE

Risk Analysis Variables Variables of risk analysis are annualized loss expectancy, annualized rate of occurrence, exposure factor, and single loss expectancy.

6. Now that the collection of facts and figures has been completed, the next step is to plug in the various calculations to determine the *annualized loss expectancy (ALE)*, which tells the analyst the maximum amount that should be spent on the countermeasure to prevent the threat from occurring. If the countermeasure costs more than the ALE, it can indicate a risk that the organization might take. This is discussed later in this chapter.

To determine the ALE, each threat undergoes the following calculation:

- 6.1. The SLE is calculated by multiplying the value of the asset by the EF:

$$\mathbf{SLE = asset\ value \times EF}$$

- 6.2. The ALE is calculated by multiplying the SLE by the ARO:

$$\mathbf{ALE = SLE \times ARO}$$

To illustrate these calculations, Table 3.2 has a short example with a few assets using a mythical Web server system.

This sample organization uses a network operations center (NOC) that cost \$500,000 to set up where the major threat is a fire. Should there be a fire, a 45% total loss is estimated. However, according to the fire department, the area where the NOC is located has a fire every 5 years, resulting in an ARO of 0.20. Using these values, the ALE for the NOC is \$45,000.

Similar calculations were made on the other assets. The asset values and EF were discovered as part of the audit; the ARO was also determined as part of the investigation. For example, when worried about power failure on the Web servers, the utility company was asked about the average length of outage in the area. In this example, the utility company predicted a major outage once every 2 years, thus resulting in a 0.50 ARO.

Using the ALE, the organization has an overview of the risks, their likelihood of happening, and what would be lost if the threat occurred. It is also known how much can be spent to protect the asset against the threats. For example, protecting against a power failure on the Web servers should cost no more than \$3,125. After some investigation, the cost of an uninterruptible power supply that works in the NOC is revealed to cost \$4,500. A business decision could be made to not employ the countermeasure because it would cost more than the loss.

-
7. The final step is to recommend countermeasures or other activities to mitigate the risk. This is the topic of the following sections.
-

TABLE 3.2**A SAMPLE CALCULATION FOR ALE**

<i>Asset</i>	<i>Threat</i>	<i>Asset Value</i>	<i>EF</i>	<i>SLE</i>	<i>ARO</i>	<i>ALE</i>
Network operations center	Fire	\$500,000	0.45	\$225,000	0.20	\$45,000
Web servers	Power failure	\$25,000	0.25	\$6,250	0.50	\$3,125
Web data	Virus	\$150,000	0.33	\$50,000	1.00	\$50,000
Customer data	Disclosure	\$250,000	0.75	\$187,500	0.66	\$123,750

Qualitative Risk Analysis

A *qualitative* risk analysis is a more subjective analysis that ranks threats, countermeasures, and their effectiveness on a scoring system rather than by assigning dollar values. There are various ways of doing this from group decisions such as the Delphi method to using surveys and interviews for their ranking system.

Doing a qualitative risk analysis is a bit different from a quantitative analysis. In a quantitative analysis, the analyst does not have to be an expert in the business of the organization or have an extensive knowledge of the systems. Using her basic knowledge, she can analyze the basic business processes and use formulas to assess value to the asset and threats. Qualitative analysts are experts in the systems and the risks being investigated. They are able to use their expertise, along with the users of the system, to give the threats appropriate ranks.

To do a qualitative risk analysis, the major threats are identified and the scenarios for the possible sources of the threat are analyzed. The scores generated in this analysis show the likelihood of the threat occurring, the potential for the severity, and the degree of loss. Additionally, the potential countermeasures are analyzed by ranking them for their effectiveness.

When the analysis is completed, the scores for the threat are compared to the countermeasures. If the scores for the countermeasure are greater than the threat, it usually means that the countermeasure will be more effective in protecting the asset. However, remember that this is a subjective analysis, so the meanings of the rankings are also open to interpretation.

Countermeasure Selection and Evaluation

Organizations employ countermeasures, or safeguards, to protect information assets. In selecting the proper countermeasures, it makes good business sense to find a countermeasure that is also the most cost-effective. Determining the most cost-effective countermeasure is called a *cost/benefit analysis*.

A cost/benefit analysis looks at the ALE, the annual cost of the safeguard, and the ALE after the countermeasure is installed to determine whether the costs show a benefit for the organization. The calculation can be written as follows:

$$\text{Value of Countermeasure} = \text{ALE (without countermeasure)} - \text{Cost (safeguard)} - \text{ALE (with countermeasure)}$$

Using the Web server example from Table 3.2, let's say that the cost of a universal power supply (UPS)—to purchase and operate—is \$1,000 per year. Even with the UPS, the exposure factor (EF) is reduced to 5% (0.05) because a power outage that lasts longer than the UPS can supply power is possible. The utility reports that an outage that will last longer than the UPS occurs once every 5 years, reducing the annual rate of occurrence (ARO) to 20% (0.20). Thus, the following calculation should be used:

$$\text{ALE (with UPS)} = \text{Cost (Web server)} \times \text{EF} \times \text{ARO}$$

$$\text{ALE (with UPS)} = \$25,000 \times \$1,250 \times 0.20$$

$$\text{ALE (with UPS)} = \$250$$

With the UPS, the ALE is now \$250. Using that for the cost/benefit analysis, you can calculate the following:

$$\text{Value of countermeasure} = \$3,125 - \$1,000 - \$250$$

$$\text{Value of countermeasure} = \$1,875$$

With the value of the countermeasure at \$1,875 and the cost at \$1,000, the benefit of \$875 per year for the countermeasure makes it a benefit for the organization.

One area skipped over was the operation cost of the UPS. The cost of operating the UPS can be a combination of power usage, modifications that might have been necessary to install the device, maintenance, and so on. When looking at the actual cost of the countermeasure during a cost/benefit analysis, all the costs need to be considered. If the countermeasure affects productivity, the loss must be accounted for. Should there be additional testing, those costs also must go into the cost of the countermeasure to get its true cost.

NOTE

Effectiveness and Functionality of Countermeasures Choosing a countermeasure for the amount of cost is a pure business way of analyzing risk. However, as security professionals, we understand that regardless of the cost, the countermeasure is not worth using unless it protects the asset. Information security professionals should work with business people to select the most effective countermeasure that will function to properly protect the asset.

This is also not a straightforward analysis. Some threats might occur once over a period of 10 years or more. Even for expensive assets, an ARO of less than 0.10 can cause the analyst to consider whether the countermeasure is worth the cost over the entire time to prevent the threat. For example, the likelihood of an earthquake destroying the network operations center in the New York City area is very low, even in an area that has seen some earthquakes. Seismologists might think that an earthquake causing some damage would occur once every 15 years (an ARO of 6.67%). But is this enough of a threat to provide countermeasures for?

Another consideration is countermeasures that can protect against multiple threats. That potential earthquake in New York might be mitigated by the rigorous building construction guidelines that keep buildings from toppling in high winds. In an information security context, a firewall can be used as a filter to prevent various network-based attacks and as a content filter to stop malicious mobile code.

REVIEW BREAK

Tying It Together

Risk assessment tells the organization what the risks are; it is up to the organization to determine how to manage the risks. Risk management is the trade-off an organization makes regarding that risk. You should remember that not every risk could be mitigated. It is the job of management to decide how that risk is handled. In basic terms, the choices are

NOTE

Residual Risk This is the value of the risk after implementing the countermeasure.

- ▶ **Do nothing**—If you do this, you must accept the risk and the potential loss if the threat occurs.
- ▶ **Reduce the risk**—You do this by implementing a countermeasure and accepting the residual risk.

- **Transfer the risk**—You do this by purchasing insurance against the damage.

These decisions can be made only after identifying the assets, analyzing the risk, and determining countermeasures. Management uses these steps to make the proper decisions based on the risks found during this process. Figure 3.3 illustrates these steps.

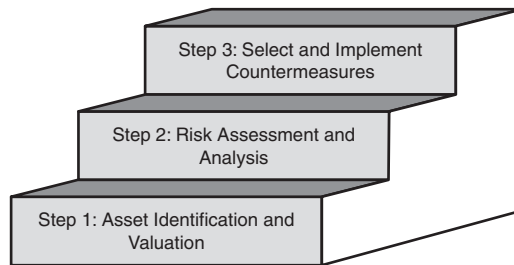


FIGURE 3.3
The three steps of a risk analysis.

POLICIES, STANDARDS, GUIDELINES, AND PROCEDURES

Know how to set policies and how to derive standards, guidelines, and implement procedures to meet policy goals.

Part of information security management is determining how security will be maintained in the organization. Management defines information security policies to describe how the organization wants to protect its information assets. After policies are outlined, standards are defined to set the mandatory rules that will be used to implement the policies. Some policies can have multiple guidelines, which are recommendations as to how the policies can be implemented. Finally, information security management, administrators, and engineers create procedures from the standards and guidelines that follow the policies. Figure 3.4 shows the relationships between these processes. The rest of this section discusses how to create these processes.




FIGURE 3.4
The relationships of the security processes.

Information Security Policies

Information security policies are high-level plans that describe the goals of the procedures. Policies are not guidelines or standards, nor are they procedures or controls. Policies describe security in general terms, not specifics. They provide the blueprints for an overall security program just as a specification defines your next product.

NOTE

Specifications Information security policies are the blueprints, or specifications, for a security program.

Questions always arise when people are told that procedures are not part of policies. Procedures are implementation details; a policy is a statement of the goals to be achieved by procedures. General terms are used to describe security policies so that the policy does not get in the way of the implementation. For example, if the policy specifies a single vendor's solution for a single sign-on, it will limit the company's ability to use an upgrade or a new product. Although your policy documents might require the documentation of your implementation, these implementation notes should not be part of your policy.

Although policies do not discuss how to implement information security, properly defining what is being protected ensures that proper control is implemented. Policies tell you what is being protected and what restrictions should be put on those controls. Although product selection and development cycles are not discussed, policies should help guide you in product selection and best practices during deployment. Implementing these guidelines should lead to a more secure environment.

How Policies Should Be Developed

Before policy documents can be written, the overall goal of the policies must be determined. Is the goal to protect the company and its interactions with its customers? Or will you protect the flow of data for the system? In any case, the first step is to determine what is being protected and why it is being protected.

Policies can be written to affect hardware, software, access, people, connections, networks, telecommunications, enforcement, and so on. Before you begin the writing process, determine which systems and processes are important to your company's mission. This will help you determine what and how many policies are necessary to complete your mission. After all, the goal here is to ensure that you consider all the possible areas in which a policy will be required.

Define What Policies Need to Be Written

Information security policies do not have to be a single document. To make it easier, policies can be made up of many documents—just like the organization of this book (rather than streams of statements, it is divided into chapters of relevant topics). So, rather than trying to write one policy document, write individual documents and call them chapters of your information security policy. By doing so, they are easier to understand, easier to distribute, and easier to provide individual training with because each policy has its own section. Smaller sections are also easier to modify and update.

How many policies should you write? I hate to answer a question with a question, but how many areas can you identify in your scope and objectives? For each system within your business scope and each subsystem within your objectives, you should define one policy document. It is okay to have a policy for email that is separate from one for Internet usage. It is not a problem to have a policy for antivirus protection and a separate policy for Internet usage. A common mistake is trying to write a policy as a single document using an outline format. Unfortunately, the result is a long, unmanageable document that might never be read, let alone gain anyone's support. Table 3.3 has a small list of the policies your organization can have.

TABLE 3.3

SAMPLE LIST OF POTENTIAL POLICIES

<i>User and Physical Policies</i>	<i>Access Control Policies</i>	<i>External Access Policies</i>
Acceptable Use	Authentication and Access Controls Encryption	Internet Security
Network Architecture	Public Key Infrastructures	VPN Access
Physical Security		Web and Internet Email

Identify What Is to Be Protected

If you remember that computers are the tools for processing the company's intellectual property, that the disks are for storing that property, and that the networks are for allowing that information to flow through the various business processes, you are well on your way to writing coherent, enforceable security policies.

The following is an example of what can be inventoried:

- ◆ Hardware
- ◆ Software
- ◆ Network equipment
- ◆ Diagnostic equipment
- ◆ Documentation
- ◆ Information assets
- ◆ Preprinted forms
- ◆ Human resource assets

It is important to have a complete inventory of the information assets supporting the business processes. The best way to create this list is to perform a risk assessment inventory. However, other methods, such as using purchase information, are available. Regardless of the methods used, you should ensure that everything is documented. Inventories, like policies, must go beyond the hardware and software. There should be a list of documentation on programs, hardware, systems, local administrative processes, and other documentation that describes any aspect of the technical business process. These documents can contain information regarding how the business works and can show areas that can be attacked. Remember, the business processes can be affected by industrial espionage as well as hackers and disgruntled employees.

Similarly, the inventory should include all preprinted forms, paper with the organization's letterhead, and other material with the organization's name used in an "official" manner. Using blank invoices and letterhead paper allows someone to impersonate a company official and use the information to steal money or even discredit the organization. So, include those supplies in the inventory so policies can be written to protect them as assets.

The most important and expensive of all resources are the human resources who operate and maintain the items inventoried. Performing an inventory of the people involved with the operations and use of the systems, data, and noncomputer resources provides insight into which policies are necessary.

Creating an inventory of people can be as simple as creating a typical organizational chart of the company. This can be cumbersome, however, if you are including a thousand, or even a few hundred, people in one document. Moreover, organizational charts are notoriously rigid and do not assume change or growth. The inventory, then, could include the type of job performed by a department, along with the level of those employees' access to the enterprise's data.

Identify from Whom It Is Being Protected

Defining access is an exercise in understanding how each system and network component is accessed. Your network might have a system to support network-based authentication and another supporting intranet-like services, but are all the systems accessed like this? How is data accessed amongst systems? By understanding how information resources are accessed, you should be able to identify on whom your policies should concentrate. Some considerations for data access are

- ◆ Authorized and unauthorized access to resources and information
- ◆ Unintended or unauthorized disclosure of information
- ◆ Enforcement procedures
- ◆ Bugs and user errors

Primarily, the focus should be on who can access resources and under what conditions. This is the type of information that can be provided during a risk analysis of the assets. The risk analysis then determines which considerations are possible for each asset. From that list, policies can then be written to justify their use.

Setting Standards

When creating policies for an established organization, there is an existing process for maintaining the security of the assets. These policies are used as drivers for the policies. For other policies in which there are no technology drivers, standards can be used to establish the analysts' mandatory mechanisms for implementing the policy.

Regardless of how the standards are established, by setting standards, policies that are difficult to implement or that affect the entire organization are guaranteed to work in your environment. Even for small organizations, if the access policies require one-time-use passwords, the standard for using a particular token device can make interoperability a relative certainty.

Creating Baselines

Baselines are used to create a minimum level of security necessary to meet policy requirements. Baselines can be configurations, architectures, or procedures that might or might not reflect the business process but that can be adapted to meet those requirements. You can use these baselines as an abstraction to develop standards.

Most baselines are specific to the system or configuration they represent, such as a configuration that allows only Web services through a firewall. However, like most baselines, this represents a minimum standard that can be changed if the business process requires it. One example is to change the configuration to allow a VPN client to access network resources.

Guidelines

Standards and baselines describe specific products, configurations, or other mechanisms to secure the systems. Sometimes security cannot be described as a standard or set as a baseline, but some guidance is necessary. These are areas where recommendations are created as guidelines to the user community as a reference to proper security. For example, your policy might require a risk analysis every year. Rather than require specific procedures to perform this audit, a guideline can specify the methodology that is to be used, leaving the audit team to work with management to fill in the details.

Setting and Implementing Procedures

The last step before implementation is creating the procedures. Procedures describe exactly how to use the standards and guidelines to implement the countermeasures that support the policy.

These procedures can be used to describe everything from the configuration of operating systems, databases, and network hardware to how to add new users, systems, and software. As was illustrated in Figure 3.4, procedures should be the last part of creating an information security program.

Procedures are written to support the implementation of the policies. Because policies change between organizations, defining which procedures must be written is impossible. For example, if your organization does not perform software development, procedures for testing and quality assurance are unnecessary. However, some types of procedures might be common amongst networked systems, including

- ◆ **Auditing**—These procedures can include what to audit, how to maintain audit logs, and the goals of what is being audited.
- ◆ **Administrative**—These procedures can be used to have a separation of duties among the people charged with operating and monitoring the systems. These procedures are where you can show that database administrators should not be watching the firewall logs.
- ◆ **Access control**—These procedures are an extension of administrative procedures that tell administrators how to configure authentication and other access control features of the various components.
- ◆ **Configuration**—These procedures cover the firewalls, routers, switches, and operating systems.
- ◆ **Incident response**—These procedures cover everything from detection to how to respond to the incident. These procedures should discuss how to involve management in the response as well as when to involve law enforcement.
- ◆ **Physical and environmental**—These procedures cover not only the air conditioning and other environmental controls in rooms where servers and other equipment are stored, but also the shielding of Ethernet cables to prevent them from being tapped.

Implementation of these procedures is the process of showing *due diligence* in maintaining the principles of the policy. Showing due diligence is important to demonstrate commitment to the policies, especially when enforcement can lead to legal proceedings.

Demonstrating commitment also shows management support for the policies. When management does not show this type of commitment, the users tend to look upon the policies as unimportant. When this happens, a disaster will eventually follow.

When enforcing the policies can lead to legal proceedings, an air of noncompliance with the policies can be used against your organization as a pattern showing selective enforcement and can question accountability. This can destroy the credibility of a case or a defense that can be far reaching—it can affect the credibility of your organization as well.

Showing due diligence can have a pervasive effect. Management supporting the administrators showing the commitment to the policies leads to the users taking information security seriously. When everyone is involved, the security posture of your organization is more secure. This does require the users to be trained in the policies and procedures, however. Therefore, training is part of the overall due diligence of maintaining the policies and should never be overlooked. To be successful, resources must be assigned to maintain a regular training program.

EXAMINING ROLES AND RESPONSIBILITY

Set information security roles and responsibilities throughout your organization.

Everyone has a role and is responsible for maintaining security in the information security process. The most important role belongs to management, who must set the tone for the entire information security program. This is not to diminish the roles of administrators and users, but without the appropriate management support, users will not take these efforts seriously.

Although information security professionals will have a more difficult time convincing users to participate in the security process, it does not absolve their responsibilities. Those whose role it is to be responsible for maintaining the information security environment should understand the roles of everyone in the organization and balance security of the information assets with the requirements of the business processes.

MANAGEMENT RESPONSIBILITY

Know what management's responsibility is in the information security environment.

Management's responsibility goes beyond the basics of support. It is not enough just to bless the information security program; management must own up to the program by becoming a part of the process. Becoming part of the process involves showing leadership in the same manner that managers show leadership in other aspects of the organization.

Management has specific goals for the organization, and most security and information system professionals are not in the position to understand or appreciate these nuances. Because security is not something that can be wrapped into a package and bought off the shelf, management must drive the attitudes for creating a good security program. This can only come after the analysis of risks, costs, and the requirements to ensure that information is not too secure to access. Management is responsible for doing the analysis and conveying this to the technical people responsible for implementing these policies.

User Information Security Responsibilities

One way to ensure that every current and future employee or user knows that security is part of his job function is to make it part of each job description. Spelling out the security function or expectations within the job description demonstrates the commitment to information security, as well as emphasizes that it is part of the job. After it is made part of the job description, it becomes something that can be considered in performance evaluations.

Outside contractors, vendors, or other people who provide external services directly on the company's network should include similar language within their statements of work. As with employees, this reinforces the company's commitment as well as makes the contractors' or vendors' adherence to the organization's security requirements a factor in their quality-of-service evaluations.

IN THE FIELD**SOCIALIZING THE ACCEPTABLE USAGE POLICY**

One common method to ensure compliance is to have anyone who accesses the network read and sign the Acceptable Usage Policy before being given access to the systems and networks. This way, users are given the opportunity to understand the policies and ask questions so they know what their expectations are.

IT Roles and Responsibilities

The information technology (IT) staff is responsible for implementing and maintaining organization-wide information security policies, standards, guidelines, and procedures. They should provide input into security awareness education programs and ensure that everyone knows her role in maintaining security. Simply, IT provides the mechanisms that support the security program outlined by the policy.

This department must be able to strike a balance between education and enforcement, although that can be difficult. They should be viewed as a partner in the business process. If implemented as an enforcement-only group, the IT group will be feared. Fear can elicit adverse reactions to their real purpose, which can undermine the purpose of these policies. Additional training can help the technology people understand their place in the environment.

Other Roles and Responsibilities

For any information security program to be successful, it must be integrated into every aspect of the environment. Integration must include statements of work and responsibilities within the business environment, job descriptions, and how these will be audited and monitored.

A primary task in assigning roles in the information security process is how information security integrates into the business environment. As part of that integration, jobs that support security through the processes should be defined. For example, one way to do this is to define a separation of duties and control over company assets by coordinating efforts with everyone, including owners of data and facilities. By having these defined as part of the business process, there is no ambiguity as to who is responsible and when.

Another role to consider is how security is administered throughout the organization. A typical environment should have a central information security management group. The central group is in charge of the monitoring and enforcement of the policy and procedures whose membership would come from the organization's stakeholders. The closer placement of security enforcement with the stakeholders can help with the control of real-time connections with third parties. These liaisons can be responsible for educating these outsiders as well as monitoring and providing enforcement.

This, however, is not a perfect solution. Some people who work in this environment for an extended period might find ways to abuse the system and exploit it, for whatever reason. One way to combat this is to not allow a person to be the security liaison for more than a short period of time—one or two years, for example. At the end of the term, they pass the job to someone else.

The final area that should have a role in the information security process is the software development cycle. Whether software is developed internally or by contractors, or if the organization purchases commercial off-the-shelf (COTS) products, the goal should be to build secure systems wherein errors or manipulations can be trapped. Policy for coding and testing standards also can assist in the quality assurance process.

UNDERSTANDING PROTECTION MECHANISMS

Understand how the various protection mechanisms are used in information security management.

Protection mechanisms are used to enforce layers of trust between security levels of a system. Particular to operating systems, trust levels are used to provide a structured way to compartmentalize data access and create a hierarchical order. These protection mechanisms are used to protect processes and data and are discussed in the following sections:

- ◆ Layering
- ◆ Abstraction
- ◆ Data Hiding
- ◆ Encryption

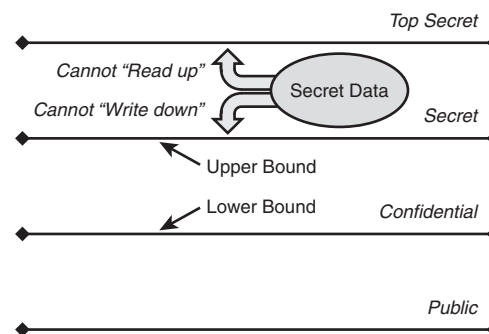
Layering

Most systems use a form of *layering* as a way to protect system resources. A traditional kernel-based operating system, such as Unix, uses a two-layer approach in which the system resources are managed in a protected kernel and everything else runs in an outer layer known as the *user's space*. If a process running in the user's space wants to access a protected resource, such as the disk, it makes a request to the kernel layer to perform the action.

Layering is specific to protecting operating system resources and to setting security zones. Systems used for military applications are designed to allow access to classified information based on the protection zone within which they are allowed to run. To do this, the Bell-LaPadula protection model was developed. Using this multilayer system, the different zones are used to keep data classified within a particular zone (see Figure 3.5). Users must have access to the zone to use the data, and the data cannot be moved between zones without special permission. This *lattice of rights* is also called “no write down” and “no read up.” See Chapter 1, “Access Control Systems and Methodology,” for more information on the Bell-LaPadula protection mode.

FIGURE 3.5

The layered zones of the Bell-LaPadula protection module.



Layering is not as common in newer operating systems. Most current operating systems rely on a set of roles and responsibilities that can simulate the layered approach. However, some specialized layered operating systems are still in use in military applications.

Abstraction

Abstraction is a common term in the world of object-oriented design. It is when data is managed as a collection called an *object*. Objects are usually defined as classes that define the data and the methods that can be used to access the object. Methods provide a predictable way to access the object's data, which allows the entire data within the class to be managed as a unit that can enforce access controls and integrity of the data.

Data Hiding

Sometimes access to data should not be provided—for example, data values within an application module that are used for internal calculations. In this case, no access methods are provided as an interface to this data. This is called *data hiding* because the data is hidden and inaccessible from the other layers.

Encryption

Cryptography is the science of creating algorithms used to encrypt data for the storage or transmission of data. *Encryption* uses those algorithms to convert data into an unintelligible form. In basic terms, encryption uses a secret key, a private value, to perform a mathematical function on the data to make it unusable by the casual observer. Traditionally, the same key is required to encrypt and decrypt the data. This is called *symmetric encryption*.

Public key cryptography is similar except that the mathematical functions can use two different but mathematically related keys. The functions generate two keys: One is kept private, and one can be given out publicly. If someone wants to send you an encrypted file, she encrypts it with your public key. Once encrypted, you can only use the private key to decrypt the message. This is called *asymmetric encryption*.

IN THE FIELD**ENCRYPTION**

Encryption is used in many areas. VPN communications are usually secured using symmetric encryption algorithms, such as the Data Encryption Standard (DES) or Triple-DES. Symmetric algorithms are used in these areas because the connections are well-defined and the exposures to the secret keys are limited.

Asymmetric encryption is used for mechanisms such as secure HTTP and email because of the multiple exposures to the keys. The public keys used in algorithms such as Secure Socket Layer (SSL) and Pretty Good Privacy (PGP) can be passed at will without worrying about compromising the encrypted channels. That can happen only if the secret key is disclosed or stolen.

Creating protection mechanisms using encryption requires several policy issues, including legal, management, and usability issues. If your organization is doing its work for the federal government, you have to consider federal standards mandated for using encryption. Encryption can be a good choice for keeping data secret, a lot of considerations must be made. For more on encryption and other cryptography issues, see Chapter 5.

CLASSIFYING DATA**Understand the considerations and criteria for classifying data.**

Throughout this chapter, we have discussed various aspects of protecting information assets. When we talk about risk analysis and management, we talk about the most cost-effective way of protecting the information asset. Part of setting the level of risk associated with data is placing it in a classification. After data is classified, a risk analysis can be used to set the most cost-effective ways of protecting that data from various attacks.

Classifying data is supposed to tell you how the data is to be protected. More sensitive data, such as human resources or customer information, can be classified in a way that shows that disclosure has a higher risk.

Information data, such as those used for marketing, would be classified at a lower risk. Data classified at a higher risk can create security and access requirements that do not exist for lower risks, which might not require much protection altogether.

Commercial Classification

Classification of commercial or nongovernment organizations does not have a set standard. The classification used is dependent on the overall sensitivity of the data and the levels of confidentiality desired. Additionally, a nongovernment organization might consider the integrity and availability of the data in its classification model.

There is no formula in creating the classification system—the system used is dependent on the data. Some organizations use two types of classification: confidential and public. For others, a higher granularity might be necessary. Table 3.4 contains a typical list of classifications that can be used for commercial organizations, from highest to lowest.

TABLE 3.4

COMMERCIAL DATA CLASSIFICATIONS FROM HIGHEST TO LOWEST

<i>Classification</i>	<i>Description</i>
Sensitive	Data that is to have the most limited access and requires a high degree of integrity. This is typically data that will do the most damage to the organization should it be disclosed.
Confidential	Data that might be less restrictive within the company but might cause damage if disclosed.
Private	Private data is usually compartmental data that might not do the company damage but must be keep private for other reasons. Human resources data is one example of data that can be classified as private.
Proprietary	Proprietary data is data that is disclosed outside the company on a limited basis or contains information that could reduce the company's competitive advantage, such as the technical specifications of a new product.
Public	Public data is the least sensitive data used by the company and would cause the least harm if disclosed. This could be anything from data used for marketing to the number of employees in the company.

NOTE**Classifications for Sensitive Data**

The classifications for the sensitivity of data used in government and military applications are top secret, secret, confidential, sensitive but unclassified, and unclassified.

Government Classification

Government classification of data is something created out of policy for maintaining national security or the privacy of citizen data. Military and intelligence organizations set their classifications on the ramifications of disclosure of the data. Civilian agencies also look to prevent unauthorized disclosure, but they also have to consider the integrity of the data.

The implementation of the classification is based on laws, policies, and executive directives that can be in conflict with each other. Agencies do their best to resolve these conflicts by altering the meaning of the standard classifications. Table 3.5 explains the types of classifications used by government civilian and military organizations.

TABLE 3.5

GOVERNMENT DATA CLASSIFICATIONS FROM HIGHEST TO LOWEST

<i>Classification</i>	<i>Description</i>
Top Secret	Disclosure of top secret data would cause severe damage to national security.
Secret	Disclosure of secret data would cause serious damage to national security. This data is considered less sensitive than data classified as top secret.
Confidential	Confidential data is usually data that is exempt from disclosure under laws such as the Freedom of Information Act but is not classified as national security data.
Sensitive But Unclassified (SBU)	SBU data is data that is not considered vital to national security, but its disclosure would do some harm. Many agencies classify data they collect from citizens as SBU. In Canada, the SBU classification is referred to as protected (A, B, C).
Unclassified	Unclassified is data that has no classification or is not sensitive.

Criteria

After the classification scheme is identified, the organization must create the criteria for setting the classification. No set guidelines exist for setting the criteria, but some considerations are as follows:

- ◆ Who should be able to access or maintain the data?
- ◆ Which laws, regulations, directives, or liability might be required in protecting the data?
- ◆ For government organizations, what would the effect on national security be if the data were disclosed?
- ◆ For nongovernment organizations, what would the level of damage be if the data was disclosed or corrupted?
- ◆ Where is the data to be stored?
- ◆ What is the value or usefulness of the data?

Creating Procedures for Classifying Data

Using this information, your organization can create a procedure for classifying data. Government organizations already have this procedure defined. Nongovernment organizations have a lot of flexibility in setting the procedures that best suit their needs. Step By Step 3.2 is an example of a procedure your organization can use.

STEP BY STEP

3.2 Creating Data Classification Procedures

1. Set the criteria for classifying the data.

2. Determine the security controls that will be associated with the classification.

3. Identify the data owner who will set the classification of the data.

continues

continued

4. Document any exceptions that might be required for the security of this data.

 5. Determine how the custody of the data can be transferred.

 6. Create criteria for declassifying information.

 7. Add this information to the security awareness and training programs so users can understand their responsibilities in handling data at various classifications.
-

EMPLOYMENT POLICIES AND PRACTICES

Determine how employment policies and practices are used to enhance information security in your organization.

Although the first concern of management might be employees and employment policies, these seem to be the last concerns of information security management. Although various research groups say that most of the threats to information assets are from internal users, employment policies can be used to protect information security assets by setting guidelines for the following:

- ◆ Background checks and security clearances
- ◆ Employment agreements and hiring and termination practices
- ◆ Setting and monitoring of job descriptions
- ◆ Enforcement of job rotation

Background Checks and Security Clearances

Those who work for the federal government, whether as an employee or a contractor, know the rigors that go into background checks and security clearances. If you work for an agency or the military

where a national security clearance is required, you probably had to fill out an extensive questionnaire that could have been verified through interviews and polygraphs. Despite some high-profile cases of personnel security lapses, the federal government does try to check everyone with access to sensitive information.

Many nongovernment organizations do not need the same type of background checks as the federal government does. However, having some type of background check should be part of the application process. Minimally, the organization should verify previous employment and other basic information provided as part of the application. For those in more sensitive positions, such as administrators and information security professionals, a further check into someone's background might be a consideration. As long as the checks are disclosed, an organization can request access to credit and criminal records to verify the applicant's suitability for her position. Organizations can even hire an outside firm that performs these checks as well as those that examine other public records to determine whether the potential for a problem or a conflict of interest exists.

Regardless of the checks your organization performs, the policies and guidelines must be disclosed to the applicant and employee. Although the government has policies for recertification security clearances, if your organization wants to do the same, that has to be disclosed to the employee. Many aspects of this are covered by federal, state, and local statutes and civil rights laws and should be cleared with an attorney before implementing.

Employment Agreements, Hiring, and Termination

In nearly every job I have had, there has been at least one employment agreement that says I will not violate policies and will maintain the integrity of the information for which I am being trusted. Other policies have included nondisclosure and intellectual property agreements. Whatever makes sense for your organization, these agreements should be presented to the new employee when he first arrives for work.

Employment agreements are used to protect the organization from something the employee can do. It is a protection from the insider threat. Agreements can also provide the organization a means by which to discipline employees if an enforcement action is necessary.

By having the employee sign the agreements, the organization has the ability to enforce the policies behind them by showing that the employee was notified of what was expected from him.

The Acceptable Usage Policy

The acceptable usage policy (AUP) is a document that summarizes the overall information security policy for the users. The AUP can contain parts of the organization's policies outlining the user's security responsibilities. Most of the time, they are highlighted components and written in plain language. A successful AUP is short and to the point. Ideally, the AUP should be only a few pages long.

Usually, the AUP is a signed document that acts as an agreement to abide by the information security policies it represents. It can be given to the new employee, contractor, or vendor with access to the network to ensure he knows his responsibilities. The purpose is to draw attention to the policy documents without requiring the new user to read them. The AUP should say that the users will abide by the policies, but the AUP can be seen as a "quick start" document to allow users to read the full policy later.

Termination

There will come a time when an employee or a contractor is no longer associated with the organization. Regardless of whether the termination is from voluntary or involuntary means, administrators must have procedures in place to revoke access to the organization's resources. Keeping a user's identification active might leave the network open for attack, and just deleting the user's information can destroy potential information assets.

Regardless of the procedures used, they should consider immediate revocation of access to the networks. Additionally, personnel policies should be adjusted to ensure employees do not have the type of access to the systems, network, and physical facilities to do damage. Even for contractors whose contracts have expired or been terminated, it might be a good idea to have a manager or security guard escort the former employee out of the building. During the process, someone should collect the employee's identification badges, keys, and other access control devices; disconnect his phone; turn off his email; lock his intranet account; and so on.

As part of the procedures, everyone must work together. If those responsible for terminating network access are not told that an employee was terminated, the network can be left open to attack by a disgruntled former employee. An improperly executed procedure makes everyone responsible for an adverse reaction.

Job Descriptions

Job descriptions are usually associated with requisitions and advertisements used to fill jobs within the organization. In the information security context, *job descriptions* define the roles and responsibilities for each employee. Within those roles and responsibilities, procedures are used to set the various access controls to ensure that the user can get access only to the resources he is allowed to access.

During periodic audits and monitoring, a user who might be accessing information beyond his job description might be an indication of a problem. For example, a contractor working on the development of the new Web system should not be able to access accounting data. The danger to this is when the job descriptions are not properly maintained. If a job description is informally changed without changing the official job description, there can be problems trying to enforce policies. It would help if there were a policy to change job descriptions before changing access control lists.

Job Rotation

Job rotation is the concept of not having one person in one position for a long period of time. The purpose is to prevent a single individual from having too much control. Allowing someone to have total control over certain assets can result in the misuse of information, the possible modification of data, and fraud. By enforcing job rotation, one person might not have the time to build the control that could place information assets at risk.

Another part of job rotation should be to require those working in sensitive areas to take their vacations. By having some of the employees leave the work place, others can step in and provide another measure of oversight. Some companies, such as financial organizations, require their employees to take their vacations during the calendar or fiscal year.

MANAGING CHANGE CONTROL

Use change control to maintain security.

NOTE**Change Control, Configuration Management, and Revision Control**

These are all similar phrases that describe the maintenance and tracking of changes to hardware and software.

The security impact of change control and configuration management is to know the present configuration of the system and its components. By knowing what is supposed to be in the system and network, administrators can identify whether security has been violated and rogue programs have been installed on the system.

One of the key security aspects of revision control and configuration management is the capability to track changes. If problems occur, administrators can examine the system in the context of the software and other installed components to see what might have caused the problem. The first step in creating these traces is to have a policy that mandates a formal change control procedure for all hardware and software systems. This policy should provide for written requests to perform system changes that can include a review for security. Using the policy as the base, the standards and procedures can be written to support the processes that log every change to any information component.

Hardware Change Control

Ideally, every time new hardware and configurations are added to the network, an entry is made to a change control system to track what has occurred. Considering that this is rarely the case, the best way to start this process is to use the risk analysis to determine the hardware inventory. With the hardware inventory, an effort should be made to place the configurations under change management control. Many organizations use the same procedures as software change management to track the changes of the configuration of the various systems. They realize that it is critical to maintain the configuration of firewalls, switches, and intrusion detection systems to ensure that someone does not change them to cover up her bad intentions.

Hardware change control does not just keep track of system and network components. Documentation should also be kept up-to-date on the network configuration, including information on where the network and telephone cables are located. Undocumented network segments might not be protected or can be used to support insider hacking capabilities. Additionally, you might want to document the various telecommunication access points into the network.

Unknown and unprotected modems can be used by anyone with access to a telephone to gain access using the software on the user's desktop, which might not be properly configured to protect the network.

Software Change Control

Software change control can have a few components. The most common topic of change control is what is used to track software development. In this case, the change management system can be used to re-create software to a certain revision to roll back from changes that might have caused security concerns or bugs.

Change control can be used to track vendor software changes. It can be considered inevitable that installed software will have bugs. Some of these bugs can be an inconvenience in operations, whereas others have security implications. It has been a source of debate among security and systems administration professionals as to how to handle fixing the software that has security problems. On one hand there is the need to fix the problem immediately to prevent problems. However, installing patches, even from a vendor, can lead to unpredicted results.

Large organizations have the capability to create test systems to test these changes before installing them into the production environment. Smaller organizations, though, might not have this luxury and might have to patch production systems. Whatever the size of your organization, having policies and procedures in place to track these changes will help you maintain the configuration of your software systems.

NOTE

Importance of Change Control

Change control on software systems can prevent unauthorized changes to those products. Untested patches can introduce bugs and other vulnerabilities that can be exploited.

SECURITY AWARENESS TRAINING

Know what is required for Security Awareness Training.

The importance of *security awareness training* and education cannot be overstated. By taking the policy, standards, and procedures and teaching all the stakeholders about their roles in maintaining the security environment, they will embrace the policy as an integral part of their jobs. This is not easy. One problem is that over the last decade,

the commitment to security by industry-leading companies has been viewed as lacking. The results are products that have insufficient security measures being installed into environments that further weaken the information security program. The dichotomy can be confusing.

Security awareness training requires clear communication. One thing you might consider for your organization is hiring a technically competent communicator for the security department. This person would do the training, educate the department to the concerns of its users, and act as a liaison between users and the department. Having someone who can communicate helps raise the confidence level users should have for the department.

Mandating that training be required for anyone with access to an organization's information assets is reasonable. Human resources should have complete records, including information on training courses required and taken as well as all signed documents showing acceptance of defined corporate policies.

Management should not only set aside time for training, but also encourage it. One company I was involved with mandated training during specific time periods, and unless employees were involved with a client or were ill, they were required to attend. This policy allowed the employee to be suspended without pay until she attended the course or watched it on videotape. You might not want to go to this extreme, but it is a good way to get 100% compliance.

CHAPTER SUMMARY

KEY TERMS

- Abstraction
- Access control
- Accountability
- Annualized loss expectancy
- Annualized rate of occurrence
- Asset valuation
- Audit
- Authentication

Understanding the management role of information security means understanding how the information security process interfaces with the rest of the organization. It is not enough to just set policies—security is a process that must be molded into the business process to support its functions. Management must support these processes with commitment and training.

Understanding what is to be protected is an important beginning of the management process. A risk analysis is used to determine the information assets that need to be protected and how they can be best protected. The risk analysis takes into consideration the costs of the assets to determine not only the countermeasures, but also whether the assets are worth protecting.

CHAPTER SUMMARY

Using this information, policies, guidelines, standards, and procedures can be created to reach the security goals. Policies can be described as the goals of the information security program. Guidelines are suggestions, and standards are the specific security mechanisms that can be used. Procedures use the guidelines and standards to implement the policies.

Access methods and protection mechanisms are used to manage the access and movement of data. A typical access method paradigm is to set the roles and responsibilities for access to the data. Protection mechanisms are used to compartmentalize access to data and processes. Layers are used to prevent unauthorized access to protected resources and data, whereas abstraction and data hiding are used to protect data.

Knowing who your users are is as important as setting their access rights to information assets. Employment policies enforce background checks during the hiring process to prevent hiring those who might be security risks. They can also set termination procedures to prevent the terminated user from destroying systems and data out of malice.

Change control and configuration management can be used to prevent unauthorized changes to the network. Change control policies can be used to maintain the configuration of all information assets to prevent them from being used to attack your organization.

The only way to really demonstrate management support of the policies and procedures is to require and support security awareness training. Through training, users come to understand their roles and responsibilities in the security environment. Training is the only way for the users to understand their responsibilities.

- Authorization
- Availability
- Awareness training
- Baselines
- Change control
- Confidentiality
- Configuration management
- Countermeasures
- Cryptographic keys
- Data classification
- Data hiding
- Encryption
- Exposure factor
- Guidelines
- Identification
- Incident response
- Integrity
- Layering
- Nonrepudiation
- Password
- Policies
- Procedures
- Responsibilities
- Revision control
- Risk analysis
- Risk management
- Roles
- Single loss expectancy
- Tokens

APPLY YOUR KNOWLEDGE

Exercises

3.1 Making Information Security Management Decisions

A good way to understand the management responsibilities of information security is to look at an aspect of a risk assessment and determine the best course of action. The following questions are designed to lead you down the decision path.

Estimated Time: 30–45 minutes

1. Your organization uses a dial-in terminal service to support customer service. The system consists of 21 inbound telephone lines and 3 outgoing lines. When calculating the risk because of an outage, the annualized loss expectancy (ALE) is \$350,000. As a countermeasure, it has been decided to look into installing another telephone circuit and modem bank. The cost for this new installation is estimated to be \$350,000, but it will lower the ALE to \$25,000. Is this a cost-effective countermeasure? Why?
2. For the previous question, which policy statement(s) should be written to support your decision?
3. Which policy statement(s) could be written that would cover the usage of the outbound modems?
4. How would you ensure that everyone knows *and* follows these policies, aside from awareness training?
2. What is the method for a system to know who is accessing its resources?
3. What is nonrepudiation?
4. What is the purpose of performing a risk analysis?
5. What are the categories of risks that are looked at during a risk analysis?
6. How are information security procedures formed?
7. The Bell-LaPadula security model uses what mechanism to protect system resources?
8. What is the difference between synchronous and asynchronous encryption technologies?
9. What is the purpose of classifying data?
10. In the context of information security, why would an organization do a background check and have an employee sign an employment agreement?

Exam Questions

Review Questions

1. What are information security's fundamental principles?
1. How do you calculate the annualized loss expectancy of a particular risk?
 - A. $SLE \times ARO$
 - B. $\text{Cost of asset} - \text{Cost of Safeguard}$
 - C. $\text{Asset value} \times EF$
 - D. $EF \times ARO$
2. What is an information security policy?
 - A. Guidelines used to define a security program
 - B. Procedures for configuring firewalls

APPLY YOUR KNOWLEDGE

- C. Management's statements outlining its security goals
- D. Risk management procedures
3. A security program is a balance of what?
- A. Risks and countermeasures
- B. Access controls and physical controls
- C. Firewalls and intrusion detection
- D. Technical and nontechnical roles
4. Which statement is true when considering the information security objectives that the military would use versus the objectives used for commercial systems?
- A. A military system requires higher security because the risks are greater.
- B. Military systems base their controls on confidentiality, whereas commercial systems are based on availability and data integrity.
- C. Only the military can make systems really secure.
- D. Military systems base their controls on availability and data integrity, whereas commercial systems are based on confidentiality.
5. What does a risk analysis show management?
- A. The amount of money that could be lost if security measures are not implemented
- B. How much a countermeasure will cost
- C. The cost benefit of implementing a countermeasure
- D. The amount of money that can be saved if security is implemented
6. Who has the responsibility to determine the classification level for information?
- A. Users
- B. Management
- C. Data owners
- D. Security administrators
7. Why should the team performing a risk analysis be formed with representatives from all departments?
- A. To ensure everyone is involved.
- B. To ensure that all the risk used in the analysis is as representative as possible.
- C. The risk analysis should be performed by an outside group and not by biased insiders.
- D. To hold those accountable for causing the risk.
8. Which of the following is not a basic principle of authentication?
- A. What the entity knows
- B. Where the entity is
- C. Who the entity is
- D. What the entity may have
9. What is the purpose of designing a system using the Bell-LaPadula model?
- A. To hide data from other layers
- B. To manage data and methods as objects
- C. To convert data to something that cannot be read
- D. To separate resources of a system into security zones

APPLY YOUR KNOWLEDGE

10. Managing an information security program is a matter of using the following principles except which one?
- A. Accountability
 - B. Integrity
 - C. Confidentiality
 - D. Availability
6. Procedures are formed from guidelines and standards to implement the stated policies. For more information, see the “Policies, Standards, Guidelines, and Procedures” section.
7. The Bell-LaPadula model uses layering to separate resources into security zones. This was discussed in the “Layering” section.
8. Synchronous encryption uses the same key to encrypt and decrypt a message. Asynchronous, or public key, encryption uses two keys: The public key of the user who is to read the message is used to encrypt that message, and the private key is used by the recipient to decrypt the message. More information can be found in the “Encryption” section.

Answers to Review Questions

1. Confidentiality, integrity, and accountability. For more information, see the section “CIA: Information Security’s Fundamental Principles.”
2. Identification and authentication is the method that associates that the object (user, process, and so on) is the entity it claims to be. See the section “Identification and Authentication” for more information.
3. Nonrepudiation is the ability to ensure that the originator of a communication or message is the true sender by guaranteeing authenticity of its digital signature. For more information, see the section “Nonrepudiation.”
4. The purpose of a risk analysis is to assess and quantify damage to information assets and to help justify appropriate safeguards. This was described in the section “Risk Management and Analysis.”
5. The risk categories are damage resulting in physical loss of an asset or the inability to access the asset, disclosure of critical information, and losses that may be permanent or temporary. This was discussed in the section “Risk Management and Analysis.”
9. Classifying data is supposed to tell you how the data is to be protected. The section “Classifying Data” explains this further.
10. Background checks and employee agreements are tools used to prevent insider attacks. This was discussed in the “Employment Policies and Practices” section.

Answers to Exam Questions

1. **A.** Answer A is the correct answer because the calculation for the annualized loss expectancy (ALE) is the single loss expectancy (SLE) times the annual rate of occurrence (ARO). Answers B and D are not correct and do not calculate anything worthwhile for a risk analysis. Answer C calculates the SLE value. See the “Asset Valuation” section for more information.

APPLY YOUR KNOWLEDGE

2. **C.** Answer C is the correct answer because policies are used to describe how an organization wants to protect information assets. Answer A is wrong because guidelines are derived from the policies. Answer B is a procedure that would support a policy. Answer D is wrong because risk management is a component in creating the policy and does not define them. See the “Policies, Standards, Guidelines, and Procedures” section for more information.
3. **D.** Answer D is correct because, as the entire chapter shows, security has both components, including physical and personnel security. Answer A is incorrect because it describes only the risk analysis process. Answer B is incorrect because it is focused on two areas of a security program. Answer C is wrong because it concentrates only on network controls.
4. **B.** Answer A is wrong because the risks can be similar and even greater for some commercial systems. Answer C is wrong because there are plenty of commercial systems that are secure, and answer D is the reverse of the correct answer. See the “Classifying Data” section for more information.
5. **A.** Answers B and C are wrong because they are parts of the risk analysis. Answer D is wrong because it is what the analysis demonstrates, which is only part of the story. See the “Risk Analysis” section for more information.
6. **C.** Answer A is wrong because the users are the ones for which the protections are being instituted. Answers B and D are wrong because they do not have the custodial responsibility to understand how data should be accessed. See the “Classifying Data” section for more information.
7. **B.** Answer A is a nice idea but not the reason to include all departments. Answer C is wrong because, even if outsiders were used, which was discussed as an option, the insiders would have to provide input into their departments’ risks. Answer D is an interesting concept, but not everyone is involved in risks. See the “Risk Analysis” section for more information.
8. **B.** Answers A, C, and D are all principles of authentication. Identifying the location can be helpful but is not one of the basic principles. See “Identification and Authentication” section for more information.
9. **D.** Answer A is wrong because it is the purpose of data hiding. Answer B is wrong because it is a principle of abstraction, and answer C is wrong because it is the principle of encryption. See “Understanding Protection Mechanisms” section for more information.
10. **A.** Answers B, C, and, D are the basic C.I.A. principles. See the “Defining Security Principles” section for more information.

APPLY YOUR KNOWLEDGE**Suggested Readings and Resources**

1. Barman, Scott. *Writing Information Security Policies*. New Riders Publishing, 2001.
2. Nichols, Randall K., and Julie J. Ryan. *Defending Your Digital Assets Against Hackers, Crackers, Spies, and Thieves*. McGraw-Hill Professional Publishing, 2000.
3. Peltier, Thomas R. *Information Security Risk Analysis*. Auerbach Publications, 2001.
4. <ftp://ftp.isi.edu/in-notes/rfc2196.txt> (RFC 2196, "Site Security Handbook").
5. <ftp://ftp.isi.edu/in-notes/rfc2504.txt> (RFC 2504, "Users' Security Handbook").
6. <ftp://ftp.isi.edu/in-notes/rfc2828.txt> (RFC 2828, "Internet Security Glossary").
7. <ftp://ftp.isi.edu/in-notes/rfc3013.txt> (RFC 3013, "Recommended Internet Service Provider Security Services and Procedures").
8. <http://csrc.nist.gov/publications/nistpubs/800-18/Planguide.pdf> (NIST SP 800-18 is a security standard used by civilian agencies).
9. <http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf> (NIST SP 800-30, "Risk Management Guide for Information Technology Systems").
10. <http://rr.sans.org> (The SANS Institute Reading Room has several individual articles that focus on many areas of information security management).
11. <http://www.rfceditor.org> (The Internet Engineering Task Force's relevant requests for comments [RFCs] are available from the RFC Editor).
12. http://www.whitehouse.gov/omb/circulars/a130/a130appendix_iii.html (OMB Circular A-130 Appendix III).

OBJECTIVES

This chapter covers Domain 4, “Applications & Systems Development Security,” one of ten domains of the Common Body of Knowledge (CBK) covered in the Certified Information Systems Security Professional Examination. We have divided this domain into several objectives for study.

Explore software/data issues and describe software and data handling applications. Demonstrate an understanding of the following:

- **Challenges of a distributed/nondistributed environment**
 - **Databases and data warehousing issues**
 - **Storage and storage systems**
 - **Knowledge-based systems**
 - **Web services and other examples of edge computing**
- ▶ If you are to understand computer security, you must understand how systems are developed. Much of what we know as computer security is the result of the inclusion of features that enable us to lock down the systems the software runs on or provide controls for data and other resource access. Most of the current problems are the result of poor development practices. Understanding the development process can assist you in understanding the scope and nature of the problems we face.

Discuss the types of attacks made on software vulnerabilities.

- ▶ What is it about software that makes it vulnerable to attacks? Is some software more vulnerable than others? You cannot explore these questions unless you know something about the types of current attacks and how they work.



CHAPTER 4

Applications and Systems Development Security

OBJECTIVES

Describe and define malicious code.

- ▶ Applications and system programs are composed of code, or instructions that can be processed by the computer. These instructions can either function normally and produce the expected results or be written to do harm. *Malicious* code is code written to do harm by making a program behave in a way it is not intended to. Understanding the types of malicious code that exist will help in understanding how to protect systems from it.

Discuss system development controls.

- ▶ It only seems logical that carefully constructed programs can be more secure programs. To carefully construct them requires strong development practices and controls. Although no one methodology has proven itself more suited to producing secure applications, understanding the major methodologies will help in your efforts to monitor and promote good practices.

Use coding practices that reduce system vulnerability.

- ▶ In addition to promoting software development controls, coding practices and good design can result in programs that are less vulnerable to attack. This section illustrates a couple of common coding and design errors that, in the past, have been the cause of major vulnerabilities that have allowed software attacks to succeed.

OUTLINE

Introduction	239
Software Applications and Issues	240
Challenges of Distributed and Nondistributed Environments	241
Nondistributed Systems	241
Distributed Systems	244
Examples of Distributed Systems	244
Massively Distributed Systems	245
Malware for Distributed Systems	246
Managing Malware	247
Database and Data Warehousing Issues	249
Data Models	251
Database Issues	252
Special Considerations for Data Warehouses and Data Marts	255
Storage and Storage Systems	256
Storage Area Networks	259
Knowledge-Based Systems	261
Developing Expert Systems	261
Techniques for Determining Answers in Rule-Based Expert Systems	261
Web Services and Other Examples of	
Edge Computing	262
Grid Computing	262
Web Services	263

OUTLINE

Attacking Software	266	Security Control Architecture	283
Attacks Against Password Databases	266	Best Practices	285
Denial-of-Service and Distributed Denial-of-Service Attacks	267	Using Coding Practices That Reduce System Vulnerability	286
Spoofing	269	Software Development Methodologies	286
Miscellaneous Attacks	270	Structured Programming	286
Illegitimate Use of Legitimate Software	272	Object-Oriented Programming	289
Network Software	273	Computer-Aided Software Engineering	291
Understanding Malicious Code	274	Impacting Security Through Good Software Design and Coding Practices	292
So, Who's a Hacker? What's Malicious Code?	275	Chapter Summary	298
Hackers, Crackers, and Phreakers	275	Apply Your Knowledge	300
Real Problems and Pseudo Attacks	276		
What Protection Does Antivirus Software Provide?	277		
Implementing System Development Controls	277		
System Development Lifecycle	278		
Waterfall	278		
Spiral Lifecycle Model	280		
Rapid Application Development	282		

STUDY STRATEGIES

- ▶ It is difficult for someone who has never written a software program or participated in a development project to understand the problems associated with developing secure programs. It is obvious, however, that something more can be done to produce software that is free from the types of errors that seem to make it vulnerable to attack. It is easy to review the types of malicious software present in today's computing environment—you've probably been in way too close contact with it. It is much harder, however, to go beyond this public view of software security. To study the development aspects of this domain requires the ability to seek the details behind the software interface to which you are exposed. Some useful approaches include
 - Study the software development methodologies presented in this chapter and review the Web links. Many times these links expose you to code examples. Reading these examples is somewhat like examining documents written with many references to foreign language examples. It's a little hard going; however, the authors often provide explanations of the code to help you understand.
 - If possible, access development documents for past projects at your company. These documents can provide you with an appreciation for the level of complexity of the software development process.
 - Obtain and read the LeBlanc (Writing Secure Code) and Viega (Building Secure Software) books on developing secure code. Although written for programmers, these books contain sufficient high-level treatments of the subject and provide interesting and understandable resources for you on software development practices that can result in a better appreciation for the degree of difficulty encountered and the beginning of a formulation for your own list of best practices.
 - Visit the sites of major PC antivirus software producers, and read the descriptions of the top ten viruses.
 - Visit the sites of security corporations and look for articles that speak to security flaws in software—that is, the why behind a vulnerability. Although sites that reveal the latest tools, exploits, and security software abound, search for those that talk about the actual code (such as www.securityfocus.com, www.eeye.com, and www.ntbugtraq.com). We all know that vulnerabilities exist. The idea is to begin to see why.
 - Organize your knowledge into the major objectives covered in this chapter, and review the terminology listed at the end of the chapter. Then, review Appendix A, "Glossary."

“Applications and systems development security refers to the controls that are included within systems and applications software and the steps used in the development. Applications refer to agents, applets, software, databases, data warehouses, and knowledge-based systems. These applications can be used in distributed or centralized environments.

The candidate should fully understand the security and controls of the systems development process, system lifecycle, application controls, change controls, data warehousing, data mining, knowledge-based systems, program interfaces, and concepts used to ensure data and application integrity, security, and availability.”

—Common Body of Knowledge study guide

This chapter covers Domain 4, Applications and Systems Development Security, 1 of 10 domains of the Common Body of Knowledge (CBK) covered in the Certified Information Systems Security Professional Examination. This domain has been divided into several objectives for study.

INTRODUCTION

On May 17, 2002, Carnegie Mellon University, Microsoft, Raytheon Co., and NASA announced the formation of the Sustainable Computing Consortium. Their goal? Write the specifications for software quality; write them so we can judge software against it; write them so that consumers will have a way to judge software; and write them so insurance companies can better judge which software or product is more likely to be hacked and thus can vary their insurance rates. Companies that use the less hackable products will get a reduction in insurance rates. Interestingly enough, the Sustainable Computing Consortium will also sport members who are lawyers, public policy experts, economists, and software engineers. You see, it's not just the “nerds” who are responsible for computer security.

Whether you consider yourself on the geeky side of this domain or hesitant to investigate it because of a predisposal to avoiding the complex subject of computer programming, you can agree, I think, to that premise. Learning about application development and the problems that can make our systems more risky gives us an appreciation for the complexity of the process and the ability to deal with excuses that point to that complexity as the reason more secure software cannot be written.

No one can guarantee that better, more secure software will be the result of your studies in these areas, but I can guarantee that your lack of knowledge of the problems and best practices will prevent your participation in what must be universal efforts to improve the quality, reliability, and security of software applications.

This chapter will help you in your studies by talking about software applications and issues, the common types of attacks made on software, malicious code, system development controls, and coding practices that can reduce system vulnerabilities.

SOFTWARE APPLICATIONS AND ISSUES

Explore software/data issues and describe software and data handling applications. Demonstrate an understanding of the following:

- **Challenges of a distributed/non-distributed environment**
- **Databases and data warehousing issues**
- **Storage systems**
- **Knowledge-based systems**
- **Web services and other examples of edge computing**

One of the problems with attempting to control the use of applications is the large range of software products that exists. It's a full day's work just to research and learn the bare minimum on the types of software and how they are used. Fortunately, you don't need to know everything. You can begin by learning about software types within the context of the issues they raise. Remember, the real issue here is the protection of the data, equipment, and lives that the software touches. Software that is unused does no harm and poses no risk. Your study should be as much about the way software is used as it is about the software itself. Begin your study by looking at

- ◆ Challenges of distributed and nondistributed environments
- ◆ Database and data warehousing issues
- ◆ Storage and storage systems

- ◆ Knowledge-based systems
- ◆ Web services and other examples of edge computing

Challenges of Distributed and Nondistributed Environments

In the beginning, there was the data center. Huge banks of metal boxes gave testimony to the fact that something was happening here. Right here; and here only. Early on, this environment shared its information only through reports. Data was entered in one location, usually by specially trained operators, and systems were maintained by other trained folks. Many of these early priests and priestesses were long-term employees selected from the ranks. New software took years to develop.

Are you getting the picture? Everything was centralized, all electronic data processing took place in one location, and it was accomplished by relatively few people. Later, this changed, and systems became distributed as new technologies and the demands of the workplace evolved. Neither centralized nor distributed systems are without their faults or challenges, and it's fitting that we begin by understanding these environments.

Nondistributed Systems

To penetrate these systems and make them run amok meant penetration of physical barriers—guards, gates, door locks, and so on. Or subversion was used—hiring on, learning the system, and then removing information or sabotaging the system. Or, the attacker could possibly coerce an employee to run a report, enter invalid data, or perform some other activity.

As these systems *grew legs*—that is, as terminals were placed in offices and directly cabled into the data center—new possibilities occurred. The terminals, “dumb” as they were, brought information to the people who used it. Information could be retrieved in minutes, sometimes seconds, and new information could be entered immediately. Although no software ran on the terminals, it didn't much seem to matter at first. Reports were still produced by the ton, and operators were still needed to punch in information from distributed locations.

NOTE

First Prominent Virus, First Problem Recognition

The first worldwide spread of a computer virus was reported in 1989. Dark Avenger, named for the author whose signature appears in the code, attached to the main operating system file—MS-DOS.com. Every 16th time the program ran, Dark Avenger deleted portions of data on the hard disk. Eventually, the computer in essence ate itself. The virus was one of 160 created in Bulgaria at the time. It spread by floppy disk and by downloading it from the first virus bulletin board.

In 1991, 600 companies were polled and 9% said they had suffered from a viral infection. Later that year another poll found 63% reported. Dark Avenger was recognized as an international epidemic. Other early viruses were Michelangelo, Jerusalem, Pakistani Brain, and Frodo. An interesting report on the early Bulgarian virus phenomena is “Heart of Darkness” at <http://www.wired.com/wired/5.11/heartof.html>.

The risks did increase, however. If everyone had the potential to directly enter data, how could you know whether what was entered was correct? Could a dumb terminal be used to attack a system? Here are some of the ways the data and the data center could be disrupted:

- ◆ Incorrect data entered in error.
- ◆ Incorrect data entered on purpose.
- ◆ Someone could enter code, which when it was run extracted data, modified data, destroyed data, and disrupted the systems operation.
- ◆ Unauthorized access to data either by getting past the controls (password sharing, password cracking, social engineering) or by seeing data displayed on screens in offices.
- ◆ Unauthorized use of unattended terminals where sessions are left active.

As you can see, the risks to software were mostly those that might be the result of bad data entry or denial-of-service attacks.

These factors also remained constant as smaller systems, the mini-computers, moved out of the data center and into the departments that used them. Accounting, finance, and marketing often justified the expense by the benefit of having local, departmental control over the data, and perhaps more importantly, the ability to ensure that their data processing projects had first priority. Systems were still isolated. Terminals were the routine, and typically, there was no interface with the corporate mainframe.

The early PCs mimicked the isolation of these nondistributed systems, when used as the main computing environment for a small business or home user. When a large number of PCs began to be used in industry, and they were networked together or other methods of data sharing were used, they became part of the *distributed computing* environment of those companies. The difference between PCs and minis or mainframes lies in their ease of use and widely distributed base. This contributed to their predilection to become attack vectors. In addition, successful software-based attacks that worked on a PC got press. In the many years of data processing when the mainframe was king, there are few records of successful software-based attacks. People didn't just decide one day to write a virus for the mainframe. What would be the purpose? Who would know?

Chapter 4 APPLICATIONS AND SYSTEMS DEVELOPMENT SECURITY 243

How many people would ever hear that there was a problem that was created by you? When mainframes and minis go down, most employees are aware only that the computer is down, not why. Contrast this to the situation in which thousands of PC systems fail due to the latest virus or worm.

PCs soon became easy victims to multiple types of software-based attacks. For these attacks to begin, the attack software had to reach the system. Like sharing sex partners, sharing data became a dangerous activity. Many malicious software programs were able to infect systems because infected files were transported between systems via floppy disks. The same types of program are still a threat in nondistributed and distributed environments. These programs fall into the following categories:

- ◆ **Viruses**—Programs loaded onto a computer without the permission of its owner and then run without permission. Several types of viruses exist, including polymorphic viruses (ones that change their own code to evade detection), boot sector viruses (those that infect the boot sector), multipartite viruses (which infect boot sectors, files, and master boot records), and macro viruses (which infect desktop application software such as Word or Excel). Often the term *virus* is used as a generic term and encompasses worms, Trojans, logic bombs, and other types of malware.
- ◆ **Trojans**—Short for *Trojan horses*, which are programs that masquerade as something else. An example is a game that, when loaded on the system, loads a virus or gathers information and writes it back to the loaned floppy disk. Another example is software that mimics the logon interface but instead captures the passwords of unsuspecting users. The perpetrator can later visit the system and, using his own, legitimate credentials, collect the captured passwords for use at a later time.
- ◆ **Logic bombs**—This software is designed to execute because of some event, such as a time (a time bomb), or as the result of some calculation or calculation result. The result can be anything from a harmless message to a system crash.

NOTE

Malware Bridges Gaps The same types of malicious programs, or *malware*, cause problems on both nondistributed systems and distributed systems. The difference is in application. Standalone systems have limited entry points, whereas distributed systems offer a broad spectrum of approach avenues (the Internet, network, media, wireless). In a distributed system they are more rapidly spread and require different techniques to thwart them and clean up after them.

Distributed Systems

As communications techniques improved and were reduced in cost, remote systems were linked to the data center via direct landline, microwave, or courier. (*Couriers* carried data in the form of punch cards, or other early data collection products, from the remote systems to data entry at the corporate headquarters and returned reports.) These were the first *distributed systems*. You should recognize the difference between distributed systems and decentralized systems. Here are some helpful ways to distinguish between them:

- ◆ **Centralized**—All computing takes place in one place. The old mainframe/data center approach is one example; another is the use of a mini-computer or mini-computers located in one place and held under the central control of one department. A single PC, used to support recordkeeping or other computing at a small company, can also be considered as centralized computing.
- ◆ **Centrally controlled computing**—In this scenario, computers can exist in a widely distributed fashion both within headquarters and at remote offices. They are, however, configured, maintained, and controlled by a central authority.
- ◆ **Decentralized**—Computing facilities exist throughout the company. They might or might not be linked with each other.
- ◆ **Distributed**—Computers are everywhere, and so is the process of processing. Distributed computing does not preclude centralized control.

Examples of Distributed Systems

It's easy to think of examples of distributed systems—we're all using them! Even the lowly home user who has no Internet connection has occasion to use computers at libraries, airports, schools, and work. These systems link themselves with many other systems by using either dedicated private networks or public networks. A few of the types of distributed systems are

- ◆ The in-house solution with PCs, PDAs, telephones, and custom devices that support entering new data to or displaying existing data from databases of information

- ◆ e-Commerce (business-to-business [B2B], consumer-to-business [C2B], and so on)
- ◆ Online banking
- ◆ Systems in which order entry occurs in one location, manufacturing in another, and warehousing in a third—all linked by some form of telecommunication
- ◆ A network of computers linked for simple file sharing or printing services
- ◆ Email

Massively Distributed Systems

Massively distributed systems are those systems that are ubiquitous across time and space and consist of hundreds or thousands of connected systems—for example, the Internet. These systems bring their own challenges. In these very large systems, people tend to trust the software more, and more importantly, they trust the results they see. This is a necessary, but interesting, paradox because users have far less control over what happens to their data after it leaves their systems. You could also say that, in order to participate in these systems, they have even given up control of their systems. They must run a browser, accept cookies, and download additional software (perhaps Java code or ActiveX controls) to fully experience the benefits their connectivity brings.

Although not the only group working in this space, The Massively Distributed Systems Group at the IBM Thomas J. Watson Research Center (<http://www.research.ibm.com/massive/>) is looking for answers to the problem of keeping computing safe in massively distributed systems. One of its projects is in developing a massively distributed immune system, one which will act to detect a new virus, develop a response to protect systems from it, and distribute a solution at a faster speed than the virus can propagate—all automatically.

NOTE

Malware; It's in the Eye of the Beholder The hosts of www.malware.org make the point that no software can be classified as malware except when coupled with the purpose behind its use. They give an excellent example of a program that formats the hard disk and reinstalls the operating system. This program, they argue, could be a useful tool when used by the administrator to prepare a new system. Yet it could be the worst type of malware if it was offered as a new game or a bug-fix program or downloaded and run without the user's permission. Their point is well taken. I think when we classify malware, we ought to at least take into account the intent of its producer and of those who offer it for use.

Malware for Distributed Systems

Examples of malware on distributed systems abound. It's just too easy to craft a tool that takes advantage of a vulnerability and then distributes the tool via the Internet or attachments in email. It's even easier if you don't have the knowledge to do so—you just download someone else's script. Information on how to create a virus for Linux (The Linux Virus Writing How-To) can be found at http://www.lwfug.org/~abartoli/virus-writing-HOWTO/_html/. You can also order Dr. Mark Ludwig's book on virus writing, titled *The Little Black Book of Email Viruses*, at <http://www.ameaglepubs.com/store/> or even obtain a CD-ROM with the source code for thousands of viruses at the same site. In addition, a site in Europe helps you write your own virus by selecting features off its Web interface.

I point these sites out, not to encourage anyone to write or to distribute malware, but to make readers aware that such tools are readily available.

What makes malware interesting to those of us who don't want to create havoc is its impact on multiple systems and the loss of productivity defending against attacks causes. Of course, we also want to defend against malware and quickly be able to clean up after it. The first step is understanding the problem.

In addition to the previously mentioned malware (virus, logic bomb, and Trojan horse) that affect the nondistributed network, distributed networks attract the following:

- ◆ **Worms**—Malware that replicates itself and spreads itself across a network. After infecting a host, the worm might use and load its own communication code, such as an SMTP engine, or simply use one of the existing services already resident on the system, such as an email, a telnet, a Web, or an FTP client. Most people have had their computers infected by some virus, and most networks have had to deal with such infections.
- ◆ **ActiveX and Java applets**—Are great examples of the intent of code helping to define whether it is malware or not. Web-based applications use these applets to do legitimate work, but malicious sites can easily pervert them and do wrong. This is a rich source for researchers. The Nimda worm relied on being able to execute Java script.

- ◆ **Blended malware**—Malware today is not limited to following the patterns recorded for it in the past. Some, such as Nimda, attacked systems that had previously been infected by Code Red. Code Red left back doors that might have not been cleaned with Code Red cleaning utilities. Nimda was also able to spread by email attachment or by download from a Web page. The sadadmin worm infected Unix Web servers and then launched an attack on Microsoft's IIS Web server.
- ◆ **Agents/remote control programs**—The capability to remotely control another computer is a useful tool for computer support personnel and administrators. However, many Trojans with remote control components pretend to be good administration remote control tools.

Managing Malware

Because malware exists in so many forms that use multiple attack vectors, no one solution will prevent its spread in a network.

Cooperation by many companies is necessary to reduce its threat to the global community. Some basic, good practices are recorded in Step By Step 4.1.

STEP BY STEP

4.1 Protecting Systems from Malware

1. Have a malware policy that specifies the use of antivirus products and provides for regular maintenance. Ensure its approval and support by top management.
2. Make virus protection software an absolute must for every server, desktop, and PDA in your network.
3. Make updating your virus protection products a priority on all systems.
4. Install and properly configure special mail server virus protection.
5. Configure mail server antivirus programs to block executable attachments.

continues

continued

6. Keep all systems patched. Many malware programs take advantage of known vulnerabilities in software.

 7. Reduce attack vectors by scanning floppy disks and other removable media before use.

 8. Reduce attack vectors by disallowing ActiveX or Java script download where possible.

 9. Keep up-to-date on trends and actual virus threats. Good practices can avoid much pain, and forewarning can also help.

 10. Use recommended steps to clean infected systems. In some cases a complete rebuild is necessary to ensure no back doors are left behind.
-

Many anti-malware products have management components with direct agents loaded on host machines. This allows efficient updating and reporting. Others propose the use of intelligent agents—code that uses rule-based inferencing engines and probabilistic decision analysis to react to malware threats. These agents might also be mobile, anti-worms, or worm cops, if you will. This is not a new idea (see J. Kephart's *A Biologically Inspired Immune System for Computers, Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, MIT Press, 1994), but it seems to be moving from the theoretical to the practical.

When I consider the prospect of having foreign code that downloads and runs on my systems to do good, I wonder whether this is such a good idea. Wouldn't it provide yet another attack vector? To learn more, take a look at the book *Mobile Agents and Security* (edited by Giovanni Vigna, 1998).

Another approach, IBM's Digital Immune System for Cyberspace, detects viral activity by using neural networks and fast pattern recognition to distinguish a virus from a nonvirus, develop a cure, and distribute it across the Internet faster than the virus spreads. Neural networks are parallel computing architectures that attempt to imitate the processing modes of the human brain. You can find an article on IBM's product—"What Is an Artificial Neural Network,"—at <http://www.emsl.pnl.gov:2080/proj/neuron/neural/what.html>.

Database and Data Warehousing Issues

Although every collection of data can be considered a database, the term is usually reserved for data that is formatted and managed by a database management system (DBMS). A DBMS consists of hardware, software, procedures, and the ability to structure the database in a way that facilitates its usage. Although there are many types of DBMS, some particular features allow their classification. DBMS defines the structure of the data and defines language syntax for accessing, storing, and manipulating the data. First, all databases usually seek to provide

- ◆ **Data independence**—Although software is provided to assist in the management of the DBMS, the software written to provide functionality for its owners does not have to be the sole user of the data. A different program can be written to use the data.
- ◆ **Minimal data redundancy**—Instead of storing data in multiple places, DBMSs make data available from multiple places.
- ◆ **Data reuse**—Data gathered for one purpose can be mined for use in another.
- ◆ **Data consistency**—Data viewed or retrieved in different ways will be the same. When a transaction is complete, the data is in a consistent state. For example, if I request the money be removed from my checking account and put into savings, after the transaction has occurred, I will still have the same total amount of money. Data consistency means it is impossible for money to be removed from one account without being placed into the other.
- ◆ **Persistence**—The state of the database and its data remains the same after code is executed.
- ◆ **Data sharing**—Many users can access the database at the same time.
- ◆ **Data recovery**—In the event of an error or a system crash, the system can recover. Transactions in process at the time of the crash are checked and either rolled back or forward to complete a transaction and maintain data consistency. The use of check points is a common technique. *Check points* are recovery points at which processing can resume after an error.

Checkpoint is also the name of a file in which the locations of the log files of the last transaction recorded to disk are listed. If a database or disk crash occurs, processing can resume at the check point—transactions in the log after this are assumed to be incomplete and are redone.

- ◆ **Security controls**—A database should be capable of providing variable security controls by limiting access to those who require it. For instance, activity can be scaled from no access to full access as appropriate to the user involved.
- ◆ **Data relationships defined by primary and foreign keys**—The *primary key* of a table is the data field or column that is used as the primary index and that allows a relationship to be built with another file. For example, the customer account number of a customer table might be identified as its primary key. Data about orders placed by this customer can be retrieved from the order file if the order file stores the customer account number of the customer who placed each order. The customer account number column in the order table is known as a *foreign key*.
- ◆ **Data integrity consisting of semantic and referential integrity**—*Semantic integrity* is enforced by rules that specify constraints. Examples of constraints are uniqueness or range matching (for example, requiring that the two initials that indicate the state in an address match one of those approved by the U.S. Postal Service). *Referential integrity* consists of the rule that no database record can refer to the primary key of a nonexistent record (if a record containing a primary key is deleted, all referenced records must be deleted).
- ◆ **Utilities or processes to ensure efficient processing overtime**—These include *compression* (the capability to compress data and save storage space and I/O), reorganization or *defragmentation* (reclaiming of unused space), and *restructuring* (the capability to add and change records, data, access controls, disk configuration, and procedural methods).

NOTE

Database Restructuring Even a database that is perfectly designed to meet the needs of today's systems might need changes as requirements change. In a relational database, for example, that might mean adding columns to a table, modifying access to columns within a table, or changing the process for backing up data. A DBMS should have the facilities for doing so in an efficient manner.

Although DBMSs in general have multiple features that attempt to ensure the security of the data, different types of databases exist. They can be classified by the way they model data. An exception to this rule is the Database Warehouse and Data Marts, which are characterized by their ability to catalog and store massive amounts of data for analysis and mining. A number of issues that can affect all DBMSs relate to the security of the data and the DBMS system itself.

Data Models

Databases are classified by the data model they use. Each model offers unique features and issues. The most common database model type today is relational, but other types of databases exist. Following are the data models that are commonly used:

- ◆ **Relational (DB2, Oracle, SQL Server)**—Data is stored in tables that consist of rows (like records in a regular file) and columns (like fields). Relationships are formed between tables based on a selected primary key. Figure 4.1 shows tables from an accounts payable database. The customer master table is related to the order table via the customer account number. The customer account number is the primary key. The invoice table includes a column that lists the customer account table. A query of the tables could easily discover the invoices related to a particular customer, as shown in Figure 4.2 for the customer Peterson's. Because Peterson's customer number is 12347, a search of the invoice table reveals two invoices.

Customer#	Name	Address	City	State
12345	ABC, Inc.	544 Smith St.	NYC	NY
12346	Johnson Tile	97 Hit St.	Atlanta	GA
12347	Peterson's	777 High Ave	LA	CA
12348	Smith & Weston	1 Main St.	Peoria	IL
12349	Bets	56 Walpole	Mexico	MO

Customer table

Primary
key

Invoice#	Customer#	Product ID	Qty	Price
567890	12347	45567	5000	1.15
567891	12349	55678	100	2003.98
567892	12347	45567	6000	1.15
567893	12348	45777	600	156.78

Invoice table

Orders Placed by Peterson's

Invoice#	Customer#	Product ID	Qty	Price
567890	12347	45567	5000	1.15
567892	12347	45567	6000	1.15

FIGURE 4.1

Defining the relationship between the customer and the order database.

FIGURE 4.2

Listing orders by customer.

- ◆ **Hierarchical (IMS)**—Data is organized in a tree structure with a tree being composed of branches, or *nodes*. Think of the branches as data records, and think of the leaves of the branches as the data.
- ◆ **Network (IDMS/R)**—Data is represented in blocks or record types. Blocks include data fields, and arrows between the blocks represent a relationship between the data.
- ◆ **Object-oriented**—Combines the object data model of object-oriented programming with DBMS.
- ◆ **Distributed**—In the typical databases (object-oriented, relational, and so on), data resides on one computer. In the distributed model, data can be partitioned across multiple computers and locations. Because the DBMS is located in many places, multiple access points exist.

Database Issues

The DBMS is designed with integrity, recovery, access control, and authorization mechanisms built in. Several of these controls must be configured or utilized. Access to the database must be granted, and granular authorizations to use the data might be possible. Backups must be scheduled and managed, and care must be taken to ensure appropriate configuration so as to not subvert any security features. Many of the security issues revolve around the database administrators' management. Administrators must understand the security features and functions of the database, be aware of security issues, and take steps to maintain them. A number of things can go wrong; here are the issues to be aware of

- ◆ **Default administrative passwords**—In older versions of SQL Server, the default SQL administrator password was blank. Many commercial products that use this database as a back end not only leave the password blank, but also will not run if it is set to anything else. Although documentation advises setting a strong system administrator (SA) password, many administrators do not. In May 2002, a new worm called the sqlsnake began circulating on the Internet. It took advantage of this vulnerability to add administrative accounts to the infected machine and send password hashes to an external mailbox.

- ◆ **Misuse, or no use, of test database**—A test database should be used for development and maintenance. After tested, code and design changes are moved to the production database. In many cases, the production database *is* the test database, or the two are connected in some fashion. This can cause a problem from two respects. First, test databases should have different administrators, which are often the programmers. However, programmers should not have access to production data because they might inadvertently modify it or expose sensitive data. In addition, a malicious employee could use this opening to steal data or corrupt the system. Second, tests and changes can make the database unstable, and data could become corrupted.
- ◆ **Lack of separation of data administration from application system development**—These duties should always be separate. The development process builds in functions that require execution by privileged users. These functions maintain the database or allow access to critical data. Database administrators, on the other hand, set access to these functions and can easily give themselves access. The programmer needs that access only during development, not on a production server.
- ◆ **Distributed databases have multiple access points**—It is hard to develop and maintain access controls across multiple access points.
- ◆ **Distributed database processing is much harder to get right**—Transaction controls, which ensure completion of a transaction or a rollback of any partial completion to the previous state, are more difficult to write for a distributed database. Typical vendor solutions provide such functionality by using special utilities, or *middleware*, to manage distributed transactions.
- ◆ **Aggregation of data can expose sensitive information**—Because of the diverse nature of the data, getting the design correct and defining appropriate permission settings are difficult. A user with no access to certain data could gain access by combining bits of data that he is allowed access to. This can sometimes be prevented by good design, but more often, it's prevented by not granting users direct access to data but rather to views. Views can be created by users with data access and can therefore be constructed to provide users with the information they need.

Yet, because the users have no direct access to data, they cannot compose a query that might expose information they should not have access to. Figure 4.3 illustrates this. In the figure, the full employee table is displayed. A box laid over the table shows the columns available from a view that has been created. Notice that the salary field is not part of the view. By providing access to the view, the database administrator has solved a privacy issue. Clerks can be given access through the view to basic employee information, but not to salary data.

FIGURE 4.3

Creating a view—access to information can be controlled.

View

Employee#	Name	Address	City	State	Phone	Dept	Title	Salary
1234	John Smith	25 Hollis Lane	LA	CA	555-1111	Sales		
1235	Nancy Willis	19 Mail St.	LA	CA	666-6547	Marketing		
1236	Peter White	444 Johnson Ln	LA	CA	555-1234	Sales		
1237	Edgar Jones	6 Butter St.	LA	CA	555-1345	Accounting		
1238	Joan Brown	555 Walnut	LA	CA	666-5678	Accounting		

- ◆ **Denial-of-service attacks**—Databases are not immune to these types of attacks. A large number of improperly formatted queries, for example, can overload the system. Examples of such queries are those that ask for complicated combinations of data tables or simply ask for every record in a very large table. Another example is a query that exclusively locks critical tables. It is normal to lock a table while essential processing occurs so that critical data does not change until the processing is done. While locked, the tables cannot be accessed by other queries. A table can also be locked by malicious code to prevent normal processing.
- ◆ **Improperly modifying data**—Data is updated from multiple sources. Unauthorized access can be gained, or authorized users can make mistakes or deliberately incorrectly modify data.

- ◆ **Access to some data can provide the ability to deduce or infer data that is protected**—This can happen if, for example, I am not allowed to access the salary records of my boss. However, I can query a view that shows the top paying titles in the company and what they pay. If I know my boss's title, I can deduce what she earns.

Special Considerations for Data Warehouses and Data Marts

A *data warehouse* is an aggregate of an organization's information. It is usually structured to provide accessible storage, query, analysis, and mining. Information placed in the warehouse is selected from that available in all areas of the organization and is typically produced by another source. Unlike the typical database, data in a data warehouse is not transactional. Instead, although new data might be appended (usually at some regular interval defined by days, not hours or minutes), most data in the warehouse is considered static and historical. It consists of large amounts of summary data collected over a long period of time. *Data marts* more typically operate at a departmental or division level.

These specialized databases can be used as a decision support system. An operational area of the company, such as sales, marketing, production, and so on, can use abstracted information to assist them in making decisions about pricing, promotions, production, and the like. Auditors might find them useful in fraud detection, compliance, and risk management. General management might find its trends capable of providing hints for asset management and cost containment. Every department might find the data rich with possibility.

So might an attacker.

The rich data sets that populate the data warehouse and empower the organization can also be goldmines for competitors, the curious, and the malfasant. Because this is not a production database, there might be a tendency to apply less security controls. Care must be taken to develop and maintain proper access controls to ensure that the data entered is correct and that only those authorized have access.

This might be slightly easier to control than for a regular database because there is less need for allowing direct access to the data by multiple people. The warehouse system can thus be firewalled off from the normal network. A thorough risk analysis should be conducted to determine where additional protection mechanisms should be employed.

NOTE

Data Mining This analysis technique requires specialized software and highly trained analysts. It looks for patterns and trends, anomalous data or activity, organized activity, and even those activities that do not follow authorized procedures.

Storage and Storage Systems

Where do data and programs live? We all know what storage is. For most of us, it's the hard drive, CD-ROM, and floppy disk we use on our own desktop systems. We've all lost changes to a document because we failed to save the file and therefore move it from RAM to drive. In the data center, larger systems use larger drives, have more RAM, and use tapes for backup. End of story.

Not really. There's more to the story, but before we descend into the details, a few definitions are in order:

- ◆ **Primary storage**—Volatile or temporary memory, it's otherwise known as random access memory (RAM). When the power is turned off or otherwise fails, any data in RAM is lost. Although more RAM can be added, there is a limit to the amount of RAM that the computer's CPU can access. The amount of RAM is also limited by the design of the computer. During the boot sequence, critical parts of the operating system are loaded into RAM and remain there until shutdown. Primary storage is faster than secondary. Data and code should remain in primary storage only while it is being used and should periodically be flushed to secondary storage.
- ◆ **Secondary storage**—Nonvolatile storage. A variety of media can store data and code for a very long time, but the media eventually decays or is replaced by other media. Secondary storage can be thought of as infinite. That is, you can keep adding another disk and move data to larger disks, tapes, or other media.
- ◆ **Real memory**—The RAM provided by the system hardware.
- ◆ **Virtual memory**—The combination of real memory and that provided by disk paging or swap files. Programs can use virtual memory addresses instead of the actual hardware real memory addresses. During operation, the program is not aware of the physical location of the data, but rather of its virtual address.
- ◆ **Sequential access**—Data is searched by beginning at the start of the media or file and searching every bit of data until the requested information is found. A typical valid use for sequential access is printing a file. Both disks and tapes can be sequentially accessed.

- ◆ **Random access**—Also known as *direct* access. Some index or other capability exists that allows a search to go directly to the record required.
- ◆ **Registers**—High-speed memory locations in the CPU. There are only a few of these locations.
- ◆ **Cache**—CPU memory storage that the CPU can access more quickly than RAM. Level 2 cache is usually a dedicated, small memory subsystem, whereas Level 1 cache is a smaller memory subsystem that is built into the CPU chip (and thus is accessible at the speed of the processor).
- ◆ **Static random access memory**—Level 2 cache that usually consists of several transistors but no capacitor.
- ◆ **Dynamic random access memory**—Memory composed of transistors and paired capacitors. Several types exist—including Fast Page mode (FPM DRAM), which processes one bit at a time, and Extended Data Out (EDO), which does not complete the processing of one bit before starting to look for and process another. Synchronous DRAM (SDRAM) uses burst mode—it stays on a row and reads ahead all data on the row. This makes it 5% faster on average because much data is read in sequence. A new type of DRAM is RAMBUS DRAM, which uses a Rambus inline memory model (RIMM) and a high-speed bus.
- ◆ **Basic input output system (BIOS)**—Provides the basic information on hardware devices, including storage devices, as well as security and boot sequences.

Consider computer memory and storage as a hierarchical structure. Figure 4.4 illustrates this. For data to be processed, it must be placed in the CPU registers. Because these are limited, some form of temporary storage is necessary. Although storage on hard disks would work, data access is slow. On modern computers, temporary storage is composed of the cache, physical RAM, and virtual storage. Storage devices are used for longer term storage and include ROM/BIOS, hard drives, removable drives, and network/internet storage (SANS).




FIGURE 4.4
The hierarchical structure of memory allows the system to efficiently work with code and data.

This hierarchical organization indicates the importance and speed of access but not necessarily the order in which it is used. During boot, critical parts of the operating system are loaded into RAM and remain there until the computer is shut down. When data is entered (either directly or by opening a file), it is first stored in RAM. When appropriate instructions are moved to the CPU's register, data can be moved to the cache for quicker access. As the CPU uses data, the data is moved back and forth between RAM and the CPU cache or registers millions of times per second. When a file is saved, changed data is stored to secondary storage but remains in RAM. If the file is closed, the memory area in which the file data exists is marked available for use.

NOTE

Locality of Reference A computer science dictum recognizes that for most programs, only small amounts of data and code are used at any one time and that often the same pieces are used repeatedly. This is why temporary memory storage works so efficiently: The same data and code are used repeatedly. You can see this principle in a different arena. More people order vanilla ice cream versus a banana split. Not only is more vanilla ice cream ordered by the store, but spare containers of vanilla are kept at the front of the freezer for easier access.

Because RAM is limited, often a portion of the hard drive is used as an extension for virtual storage. Instead of placing data back in a permanent location on the hard disk, data is temporarily placed in a paging file and can more rapidly be located and moved back to RAM as necessary. You should be aware that this file might not be cleared at shutdown. Although it is protected from direct access by anyone other than the operating system, while the computer is operational, it is logically represented as a file on the disk. Should an attacker gain physical access to the computer, he could boot it to another OS and make a copy of the paging file. He then could analyze it and potentially find sensitive data. In some operating systems, you can schedule the paging file for clearing at shutdown.

The following lists the storage devices and the types of memory they represent:

- ◆ **Credit card memory**—A special, proprietary, DRAM memory module that can be used by placing it in a slot on a notebook computer.
- ◆ **PCMCIA Card**—A nonproprietary DRAM module that works with notebook computers designed to the standard.
- ◆ **Flash RAM**—A small amount of refreshable memory used by cars, TV sets, VCRs, and so on to remember configuration data. Even with the power turned off, the chip can access a small amount of power to keep itself refreshed. It is often used on computers to store hard disk information.
- ◆ **Real-time clock (RTC)**—An onboard chip on PCs that keeps time. The 64-bit of RAM also stores floppy and hard drive configuration information needed during boot. This RAM is kept alive by a small battery, called the CMOS battery, even when the computer is turned off.

- ◆ **Video RAM (VRAM)**—Also known as multiport dynamic random access memory (MPDRAM), it is used for video adapters and 3D accelerators.

Storage Area Networks

Storage area networks (SANs) are centrally managed network accessible storage systems. In a typical environment, they are accessible from all servers and other storage systems. Their benefits are many:

- ◆ Centralized control, including backup and management.
- ◆ Access from anywhere at anytime.
- ◆ Can improve data protection.
- ◆ Additional storage can be added with little to no disruption.
- ◆ Better physical security.
- ◆ Improved availability.
- ◆ Business flexibility.
- ◆ Can improve disaster tolerance.

The first SANs introduced relied on the Fibre Channel protocol—not a typical attack vector at the time. The obvious concern here is that, as knowledge of SANs usage and architecture grows, so will the attention of those more likely to be eager to attack them. Historically, if the fruit is juicy enough, someone will find a way to obtain it. Obscurity never suffices as security for very long.

Now that SANs are moving to the use of IP-based networks, they will be vulnerable to the attacks presently deployed against other services on these networks.

To secure SANs networks, insist that SANs products have the features that will allow SANs administrators to apply these general security principles:

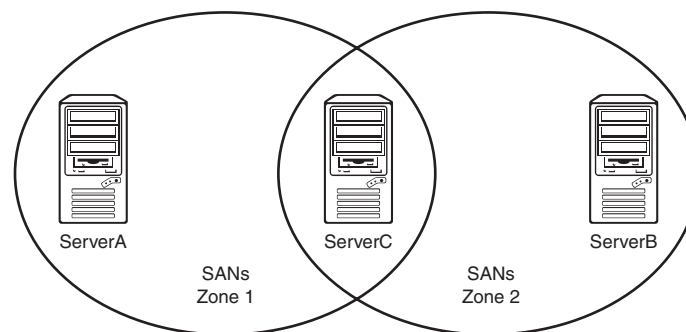
- ◆ **Physical security**—Where SANs devices are centrally deployed, this is an easier task if they are contained in secure data centers. However, distributed SANs—those SANs with devices at remote locations—will be harder to secure.

NOTE

SANs Technology Reference There is no lack of information on SANs. Every vendor provides a wealth of data on its site. IBM, in addition to general information and product-specific documentation, provides an online viewable redbook called *Introduction to Storage Area Networks, SANs* at <http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg245470.html?open>. This book also provides an introduction to the SANs standards organizations and standards.

FIGURE 4.5

Creating SANS zones allows the maintenance of access rights when new SANS are added and therefore can assist in securing data.



- ◆ **Confidentiality**—If SANs will use IP networks, encrypt SANs data in transit; IPSec can be used to do so. This will not prevent IP sniffers from capturing data, but it will prevent them from reading it. A SANS also can support local encryption of the data to secure it during storage.
- ◆ **Authenticate users**—All access to SANs data should rely on established mechanisms for validating the identity of the individual.
- ◆ **Authorization**—Access controls should have granular application. Setting appropriate access to data is a mandatory feature of any storage system. Although application-level controls are important, other mechanisms should be available. The typical file and folder access controls available with modern operating systems are useful additions to a SAN. An additional control available in some SANs is the ability to zone, or segment, SANs devices and make only some devices available from some servers. Figure 4.5 shows this technique. In the drawing, ServerA can access SANs devices in Zone 1. ServerB can access only devices in Zone 2, and ServerC can access devices in either Zone 1 or Zone 2. As the need for more storage grows, devices can be added to the appropriate zones without changing the access rights of any of the servers. This can be accomplished with IP switching. (You should survey the literature on the viability of this technique. IP switching is not a security mechanism.) Another fallacy is also embodied here: What if an attacker obtains control of ServerA? She would then have unlimited access to the SANs devices in Zone 1.

- ◆ **Interoperability**—When different vendors' SANs are used, difficulties can exist in communications between them. This can cause security problems because security controls in one SANs might be reduced to accommodate the lack of security controls in another.

Knowledge-Based Systems

Knowledge-based systems, often called *expert systems*, attempt to parallel the thought process and deduction efforts that transpire when an expert searches for the answer to a problem. In one model, the expert examines known data and asks a series of questions whose answers lead to more questions until the answer is found.

For example, take the common workplace question, “Where should we go to lunch today?” As the local expert, you know many places that serve lunch. You might begin the process by asking what type of food the others want to eat. If Sally says, “Anywhere but nowhere expensive,” you immediately reject your favorite restaurant, Chez Topos. If a consensus is finally reached that includes Mexican or Italian, you react by filtering your list for only Mexican or Italian restaurants. Next, you ask about transportation and find out that no one drove today; thus, you reduce the list to the only restaurant in walking distance.

Expert systems use a similar technique to solve problems. They use a set of rules against known data to infer new information.

Developing Expert Systems

To develop such a system, you use an expert system shell, which consists of an inference engine and a user interface. The developers add the data in a specialized format and write the rules. Often the data and rules are developed during consultation with experts in the field that the expert system will exemplify. This process, the taking of expert knowledge and codifying it in a database, is known as *knowledge engineering*.

Techniques for Determining Answers in Rule-Based Expert Systems

A rule-based expert system, such as those described previously, is populated with a database of information and a series of if/then rules. The answer to a question is found by one of two techniques:

- ◆ **Forward chaining**—This begins with a question and a set of known facts and proceeds to evaluate related rules. If a rule is true, it fires and produces more information, and thus more rules can be evaluated. The process ends when no new facts can be obtained or the result for the question is found.

NOTE

Seeing Is Believing Sometimes it's hard to translate a definition on paper into something the mind can relate to. Working knowledge-based systems are present in the real world; seeing the underlying processing that makes them work is often difficult. You can visit <http://www.emsl.pnl.gov:2080/proj/neuron/kbs/demos.html>, which provides links to research projects that demonstrate knowledge-based systems, often with information, flashing lights, or other devices that help you understand the event firing or other processing.

Another site of interest is <http://www.expertise2go.com/webesie/tutorials/ESIntro/>, which demos an expert system and introduces terminology along the way.

Finally, the existence of real systems in use today is always a good reality check. You can find medical expert systems listed at <http://www.computer.privateweb.at/judith/> and a story about CYC, a computer loaded with common sense and that is now a product called CycSecure. CycSecure knows what a hacker can do, knows what normal activity is on a network, and can be used to logically test your network's defenses by deduction, not by actual hacking (<http://www.cnn.com/2002/TECH/industry/04/11/memome.project.idg/index.html>).

NOTE

Grid Computing Resources To learn more about grid computing, check out these resources:

- www.globus.org—Access the Globus projects site, where you will find articles, research information, current projects, and even tool kits if you're interested in developing grid computing in your organization.
- www.beowulf.org—For an example of Linux clustering, see the Beowulf project site.
- <http://oscar.sourceforge.net/>—For an open source project, visit this site.
- http://www.gridcomputingplanet.com/news/article/0,,3281_1365171,00.html—For an example of a grid see this site, which tells the story of Purdue University and Indiana University linking their supercomputers to form a teragrid (more than a teraflop of processing power, or the ability to process more than a trillion operations per second). Its purpose? Simulating terrorist attacks to help government agencies plan ways to mitigate the effects of such attacks, or at least deal with the aftermath.

- ◆ **Backward chaining**—This starts with a hypothesis or question that can determine the answer and then works backward through the rules attempting to determine whether the answer is correct.

Web Services and Other Examples of Edge Computing

Throughout the history of computing, new technologies have arisen because of a need or because the hardware had finally caught up to the imagination. Interestingly enough, current advances are centered on the premise to push processing from large, centralized data centers to distributed foci—rather than requiring bigger, more powerful computers. Grid computing allows the gathering in of the excess processing capability from the proliferation of computers in the typical organization—it does for processing what SANs are doing for storage. In contrast, Web services dissect the program processing into its smallest chunks and spread the program's pieces across the Internet, thus allowing these chunks to be recombined in many different ways.

Grid Computing

If all processing power is located in a single computer, the computer can be designed to take advantage of idle moments to run less critical programs. In a modern organization, processing power is spread across thousands of servers, mini-computers, and desktop systems. Although each computer serves a purpose, many of them are not used to capacity. Think of the average desktop computer. Is it used all day long at full capacity? How about at night and on the weekends? What if the idle processing time of all the computers in the enterprise could be harnessed and utilized? Grid computing seeks to do this.

The concept of combining multiple processors to solve complex problems quickly is not a new one. You might remember the Cray computer (first created in 1976), early versions of which were used in weather forecasting and various simulation projects. Multiprocessor computers mean big bucks. Although these systems are cheaper now, they still require significant outlay. *Clustering*, or the combining of multiple computers for the sharing of processing power and storage, is a more recent development.

Indeed, there are those who define grid computing by pointing to the cluster. But even with the cluster, the computers are dedicated to the tasks assigned to them.

Enter grid computing, a structure in which excess processing power is made accessible for new tasks. This harnessing of idle time can be accomplished within an organization or work across boundaries connecting disparate machines across the Internet. Envision, if you will, a future in which you can sell the excess capacity of your computers much like power companies broker excess kilowatts. Grid computing also means the capability of software to aggregate other computer resources, such as information. In some ways, Microsoft's .NET is a computing grid that distributes processing over multiple computers.

An interesting article on grid computing is "The Anatomy of the Grid, Enabling Scalable Virtual Organizations," written by Ian Foster, Carl Kesselman, and Steven Teucke and published at the Globus Web site (<http://www.globus.org/research/papers/anatomy.pdf>).

You can participate in a grid computing project; in fact, you might inadvertently be doing so. Some "free" services or downloads come with software EULAs (licensing agreements) that authorize the parent company to use excess bandwidth or processing power in your network! Grid computing projects also exist, such as SETI (<http://setiathome.ssl.berkeley.edu/>), which seeks to harness excess cycles on home computers to facilitate extraterrestrial research.

Web Services

What do you use the Internet for? Do you use it to send and receive email? Bid at an auction? Purchase books, clothes, airplane tickets, or other things? Research information to help you in your work? Many services are available on the Internet. Some of them are available to the public, and others are open only to registered users or represent private transactions between divisions of a company or between companies. But these services, though useful, are not necessarily "Web services." One definition of Web services is that they are small, reusable programs that can be accessed from otherwise unconnected sources. Web services can be written in XML and used to communicate across the Internet or an organization's intranet.

NOTE

Can the Holodeck Be on the Horizon? On the starship Enterprise, an empty room becomes whatever you want it to be. Simulated people and complete worlds exist where participants can enjoy a vacation, solve a mystery, or picnic with a long lost loved one. When Purdue University and Indiana University combine their supercomputers, they'll be able to simulate the actions of real people, in hopes of solving real-world problems. Although no plans for a *Star Trek*-style holodeck for entertainment are revealed in their publicly listed project scope (and even this megapowered grid probably is not capable of putting bodies and worlds together for real people to walk through) can such designs be far behind? And if so, who would create them? A new report, "Global Grid Computing Report 2002: Technology and Market Opportunity Assessment," by Grid Technology Partners (www.gridpartners.com) gives an example of how grid computing can bring more power to companies: "A company with 600 grid-enabled desktop PCs can utilize all of them together as one computer platform—suddenly providing it with enough computing capacity to go head to head with the world's 49th largest supercomputer" (http://itmanagement.earthweb.com/it_res/article/0,,3031_1033451,00.html).

A traditional piece of software incorporates all the code it needs within a program. Although *code libraries* (collections of reusable functions such as DLLs) can be used by more than one program and modules can be distributed across multiple platforms, the fact is that control and management of the entire body of code are under the control of a single program. Various approaches have been used for the sharing of code across systems. Remote procedure calls, DCOM, and CORBA are examples of the way this is done. The use of these techniques to process data between companies is hard because each company must struggle with interoperability, reliability, and security problems. On an individual basis, and between companies, the Internet protocols HTTP and HTTPS are often used and programs that allow the transfer of information between two different organizations are created. Although this can work, for many the inefficiencies are rampant because applications must be entirely rewritten.

Web services can solve these problems by using XML, the universal language for data exchange. Although Web services are far from being fully developed, many companies are using this technique to develop new applications and even to wrap legacy code. Web services can work in many scenarios, including

- ◆ **Client-to-client**—Web services can share data between clients.
- ◆ **Client-to-server**—The traditional “me boss, you slave” orientation to data collection, analysis, and recall.
- ◆ **Server-to-server**—Processing can take place across multiple servers—anywhere.
- ◆ **Service-to-service**—Services can work together, in Web services.

It is tempting to think of Web services as just another programming paradigm that varies only slightly from every other already available. The difference is this: If you can think of a typical program as being a collection of small steps and a modern program as an efficient arrangement of these steps into subprograms or functions, then you can think of modern distributed processing as some master program that periodically accesses some subprogram which exists on some device somewhere. Taken a step further, many programs work in their own space, often accessing bits of code resident elsewhere and occasionally communicating with each other to share data.

Now, take a gigantic step: Think of each subprogram as existing entirely independent of every other subprogram. Processing occurs when these subprocesses are strung together like so many scenes from good movies. In short, we've moved to collage computing.

Here's an example: Today, if I want to schedule a flight to Seattle, I can book it online either through an airline's site or at one of the aggregation sites, such as Expedia or Orbitz. When I do so, my request for available flights is processed entirely by the site I'm connected to. It might access other sites to compile information and present it to me, and it might recontact those sites to actually book my flight after I've purchased it from the site. In short, it acts as a travel agent, gathering and then feeding me information I request and making arrangements for me after I decide. It might even ask whether I also need hotel or car rental services, but it is the Web site which is the aggregator.

In the future, Web services at each airline will advertise available flights and rates. At hotels and car rental agencies they will do likewise for their services. Web services at the aggregator, instead of complex applications, will work with the Web services of the other companies to obtain data that they then merely feed into their proprietary formats. There might even be a Web service resident on your computer that can independently visit multiple airlines and compile composite information.

In the past, much work on the part of the aggregator and the airline was necessary to build communication links and process between them. With Web services, the airline could build the Web service once and any aggregator running Web services (perhaps a plug-in to my browser) could access them.

Web services can also solve the problem of interim information. Because my travel will be on an airline and my contract is with an aggregator, what happens when the airline adjusts its schedule? Currently, the airline notifies the aggregator who, hopefully, notifies me. The aggregator does not want me to be directly contacted by the airline—it might lose me as a customer. With Web services, my resident Web service might, with information obtained from the initial transaction, periodically query the airline Web service for updates. When an update is received, some form of alert might be communicated to me (possibly on my cell phone or PDA).

NOTE

Reality Check Web services tool kits for Microsoft Office XP let users pull data into Excel spreadsheets from Web sites. To do so, the Web sites must host Web services. The tool kits help develop them. Companies that have used the tool kit include FedEx, Jet Blue Airways, and General Motors. Maybe my previous example of the user aggregating the data is not so far behind. To find out more about Microsoft Web services, visit the following: <http://www.microsoft.com/net/defined/whatis.asp>.

NOTE

A Thousand Points of Light What president said that? (It was George Bush.) He was speaking of creating a nurturing climate for education and envisioning new efforts at schools as shining points that spread across the country. If you close your eyes, can you see Web services as small lights spread across the Internet?

ATTACKING SOFTWARE

Discuss the types of attacks made on software vulnerabilities.

To write or select good software and to protect it from compromise, you must understand how software is developed, the controls that are available during its production, and the types of attacks that are directed at software. This section enumerates on the latter.

Many attacks on software are based on flaws, whereas others are directed at the inherent weaknesses in the components, protocols, and processes from which software is built. Still others work by subverting the process and placing malicious code within an otherwise innocuous application. The following sections discuss the typical attack types that are often utilized.

Attacks Against Password Databases

Two common types of attacks against password databases are *brute force* and *dictionary* attacks. These are attacks against the use of weak passwords. They could be considered attacks that are made possible due to weaknesses in the password policy facilities of the OS, or due to weaknesses in the authentication protocol used by the OS. For an example of the former we have only to look at early versions of Unix, which placed unencrypted passwords in a file and attempted to use obscurity and file access permissions to keep the file contents safe. An example of the latter is the use of the Lan Manager (LM) authentication protocol by Microsoft Windows 95 and Windows 98. This protocol has several well-documented weaknesses that make a brute force attack easier to accomplish. This protocol is not used by more recent Microsoft OSs. A brute-force attack seeks to determine a password by trying every possible combination of characters.

A dictionary attack is successful because many users will use regular words as, or as part of, their password. The attack encrypts common dictionary words with the same algorithm used to encrypt passwords and then compares the encrypted passwords to the password file. A match, of course, means the password has been discovered. A more sophisticated tool will also determine whether regular dictionary words are part of the password. Typically, the tool comes with a dictionary file but allows for adding words or entire dictionary files.

In this manner, company-specific words or larger dictionaries can be added.

Adding additional authentication factors such as smart cards, biometrics, or other devices can prevent these types of attacks from being successful; however, they do offer a different attack surface, albeit a harder one to crack. If passwords are used, make it a rule to use strong passwords and to implement a strong password policy. This can only be accomplished if both operating system facilities are used and user cooperation is obtained. Longer passwords—composed of uppercase and lowercase letters, numerals, and punctuation—frequent password changes, and prevention of password reuse are some of the techniques that can be used. It is also essential to train users in the policy and in how to create strong passwords that are free of words that might be in a dictionary. Operating system vendors should assist by including the facilities to make stronger password policies and including software that enables the use of other authentication factors.

Many commercially available and freeware password crackers use this technique. Jack the Ripper and LC4 are examples.

Denial-of-Service and Distributed Denial-of-Service Attacks

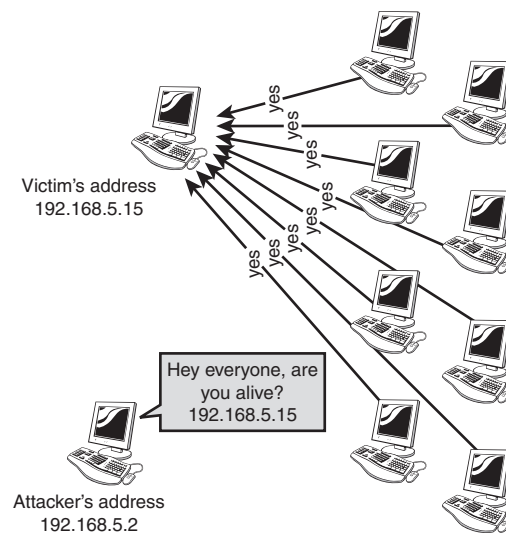
Many forms of denial-of-service (DoS) exist. In fact, many types of attacks result in a DoS. Others might have that as their effect, but might also offer more serious consequences depending on the security context of the exploit—for example, if the attack is executed by, say, tricking a user to run it. If the user is unprivileged, the security context in which the exploit runs is unprivileged, and perhaps little will happen. If, however, the user is a systems administrator with root privilege, the exploit runs within this security context and more serious consequences result, including perhaps adding new privileged users.

Although denial of service can be the inadvertent result of many attacks, some attacks have this as their purpose. A DoS, though, is any attack that is successful in keeping legitimate users from the services the computer software offers. This might mean crashing the system or the software, merely tying up connections to the computer, or accessing the software or its database in such a way that no legitimate user can gain access.

For example, the attacker might crash the server by overflowing the buffer of some data entry point. Much code is written that does not check the length of data entered by the user. When long strings are sent, instead of the expected information, a system crash or worse can be the result. For more information, see the section “Eliminating Buffer Overflows.”

Another DoS, called a *smurf* attack, is the result of sending a spoofed source address in an ICMP ping packet to the broadcast address, thus causing all computers on the network to send a response to the victim (at the spoofed source address). The ICMP ping command seeks to see whether a computer can be located on the network. When it is used, the source address—that is, the IP address of the computer that is used to issue the command—is automatically entered into the packet that traverses the network. The destination address is the IP address of the computer that is sought. If that computer is on the network and receives the request, it returns an answer to the source address. However, an attacker might craft a packet and place the IP address of his victim as the source address. If this packet is sent as a broadcast (meaning, it would be received by every computer on the network), all computers would answer by sending a response to the victim. This might overwhelm the victim, hence causing the DoS. Figure 4.6 illustrates the problem. The solution is software that prevents such a problem and indeed, most modern TCP/IP stacks are so written. This attack is one that can be successful if there is a software flaw.

FIGURE 4.6
The classic smurf attack.



A distributed denial-of-services (DDoS) is a DoS that is accomplished by first gaining control of multiple computers and then using them to attack the victim. Figure 4.7 illustrates the concept.

A reference on DDoS can be found at <http://staff.washington.edu/dittrich/misc/ddos/>. Here, links to articles on examples of DDoS, such as Trinoo and Tribal Flood network, as well as other material can be found.

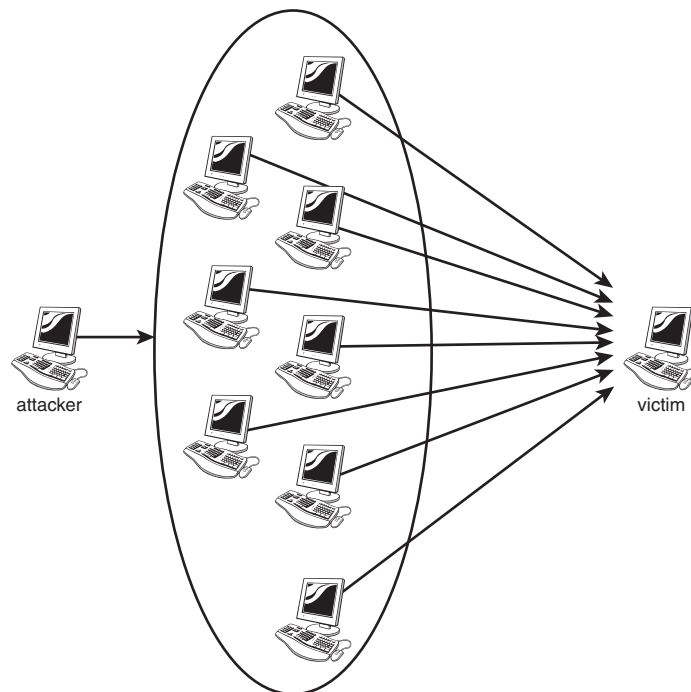


FIGURE 4.7

Distributed denial-of-service attack. In the diagram, the attacker is controlling multiple PCs or zombies to attack another PC, the victim.

Protection against many forms of DoS consists of the application of all current patches and service packs. For other types of DoS, the solution will only come when all software is written to prevent buffer overflows. Still other attacks cannot be prevented except by blocking traffic from the attacker. DDoS attacks will be possible as long as there are vulnerable machines on the Internet.

Spoofting

There are many types of spoofing attacks and many attacks use some form of spoofing to accomplish their goal. We have already discussed one, the smurf attack, in which an IP address is spoofed.

NOTE

Flooding In mid-2002 a new worm began to move across the Internet. It sought to take advantage of a software flaw in the Apache Web server for FreeBSD in order to make the server a zombie. You're correct if you equate that with the mindless creatures under the control of the evil monster in some twentieth-century horror flick. Computer zombies are under the control of a master. The worm was trying to create its own stable of compromised machines, a flooding net, that it could then use in a coordinated attack against some new victim.

Spoofing, then is the attempt to use the credentials of another computer in order to accomplish some goal. Several different techniques are used, and different credentials can be spoofed:

- ◆ To simply direct an attack at a victim (the smurf attack).
- ◆ To gain entry to a network where the MAC address (the address of the network card), IP address, or name of the computer is used for authentication.
- ◆ To take the identity of a host computer in order to act as that computer in some man-in-the-middle or similar attack.

The SMBRelay attack is one example of a spoofing attack in which an attacker attempts to take the identity of a trusted host by using the MAC address. SMBRelay is a tool that hijacks a Server Message Block (the communications protocol used for Windows file sharing) session between two computers. SMB signing, a process that authenticates each packet in the file sharing session, can be used to prevent the success of the SMBRelay attack.

Miscellaneous Attacks

Software-based vulnerabilities include intentional misrepresentations, accidental inclusions, and poor design. Examples of each of these are as follows:

- ◆ **Hidden code**—Code can be inserted within an approved software program. In poorly managed code, where code review is not done, this can be easily accomplished by a member of the team. Otherwise, special techniques might be used. One technique uses the NT File System (NTFS) or other file systems that use file streams. This is a little known capability of NTFS and it quite easily could be used to hide code. Although it is easy to view the code if you know it is there, finding which files might be using file streams is not an easy task. Another technique would be to develop and use a virus to hide code within existing code. Viruses typically attach themselves to existing code so that they can hide. Vet, or approve as trustworthy, application development teams and audit their work. Scan code for the use of file streams, viruses, and such.

- ◆ **Logic bomb**—A program that lies dormant until activated by some event, say a time, or by the use of a specific program. Often these are placed on a computer by a virus, but more often are the result of some disgruntled employee hiding code within an approved program. The bomb is set to “explode” or go into effect when the employee (or former employee) is not present. Audit activity involving code maintenance, code production, and access to servers. To find logic bombs, use virus-checking programs. (These programs will work for known viral code with delayed “logic bomb” action. A virus checker will not protect against the activity of an employee writing his own code.)
- ◆ **Trap door**—During program development, access to operating system debugging facilities is often programmed in at specific points as a programmer aide (program debugging break points). When the program is moved into production use, or offered for commercial sale, these “trap doors” or portals that circumvent system protection, should be removed. Some trap doors can be activated by typing a set of keys. The idea here is similar to a back door, or a way to gain unauthorized entry to a system. If the programmer-debugging tools are not removed, they might be used to compromise the system. The existence of break points still present in production code can be the result of deliberate lack of removal or as the result of careless development practices. Trap doors can also be accidentally created by the combination of unintended combinations of functions. To prevent possible compromise due to trap doors, insist on code review and look for removal of break points and other programmer-debugging techniques as well as unusual code. Functions should also be tested in all combinations.
- ◆ **Time of Check to Time of Use (TOC/TOU)**—If an instruction is executed in more than one step, it might be possible to compromise the system by attacking between the steps. TOC/TOU is the name for a special type of race condition that can be vulnerable to this type of attack. IBM’s OS 360 (an older mainframe system) performed access control over files by first reading and checking permissions; then, if the permissions were correct, the file would be read again. If the permissions were incorrect, the user would be denied access.

However, if the system could be interrupted before the denial was returned, the file could be read and possibly modified. More recent race conditions (conditions that exist because of timing issues within software) include problems with the `rm` command in Linux. Because of the way the command was written, it could be reissued before complete, causing a DoS for an unprivileged user, and a possible removal of the entire file system if the user was a root user. This error is not present in updated versions of the OS. You can read more about it at http://www.linuxsecurity.com/advisories/caldera_advisory-2045.html.

- ◆ **NAK attacks or interrupts**—Interrupts are used by devices to alert the operating system to their need for attention. Examples include a key press on a keyboard, and the arrival of data at a modem port. Software interrupts are also used. When a service is requested, the typical response is with an acknowledgement, an ACK, or a negative acknowledgement, a NAK. If a system is not programmed to properly handle these interrupts, the system might be left in an exposed state. A NAK attack takes advantage of this.
- ◆ **Pseudoflaw**—Have you ever tried to do something on the Internet and had it fail? You were told to try again. When you did, did you succeed? You might have been the victim of a pseudoflaw. This type of attack might insert its own code in front of or around the real code. In a logon pseudoflaw, the victim enters her user ID and password and is told she has entered an incorrect user ID or password. When she tries again, she succeeds. She might think she simply mistyped, but in reality the pseudoflaw recorded her user ID and password and then returned her to the legitimate logon screen.

Illegitimate Use of Legitimate Software

Administrative tools, in the hands of the wrong person, can be as destructive as any hacker tool. Sometimes, a tool that pretends to be a legitimate tool is really a hacker tool in disguise.

Legitimate remote-maintenance programs are a necessity in modern computing; there are just too many servers to manage, and too many desktops to service to be attending to them individually from the console. Unfortunately, legitimate programs can be used by malicious users, and rogue management programs also exist. SNMP is used by many remote-management programs. Recently exposed vulnerabilities in SNMP, if not patched, make it useful to attackers, who can use it to control systems, or to learn information about them.

Netbus, and Back Orifice are remote-control Trojans that enable infected machines to be controlled by another machine running the matching client. The Trojan “server” can infect the victim machine by tricking the user to load some other program or clicking an attachment. There even are administrators who think the program is okay to use for remote administration. They don’t realize that the tool might also have embedded software that makes it easy for an attacker to locate and control those machines the administrator thought he had protected from unauthorized management.

Netcat, though used by some as a legitimate network management tool, can also be used as a Trojan. If an administrator loads netcat on a PC and schedules it to run and listen on port 23, the administrator can obtain a command session on that PC and thus manage it remotely. However, because no authentication takes place, an unscrupulous attacker can also command the machine. In addition, if I tricked you into running netcat on your system in such a mode, I’ve in effect trojaned your box. Once again, interpretation of the use of a program labels it either a “network management tool” or a “Trojan.” Visit <http://www.atstake.com/research/tools/nc11nt.txt> for a description of some of the useful things that can be done with netcat—by the network administrator or an attacker.

Network Software

A server can be vulnerable due to flaws in the software or it can be at risk simply due to the role it plays. Likewise, the networking software and hardware that connects the computers on your network might also put it at risk. Examples of this are plentiful; here are a few:

- ◆ In a Windows network with browsing enabled, computers show up in the browsing window. When clicked on, those computers reveal shares, or entry points, to the hard drives.

If permissions are set properly, these folders cannot be accessed by the unauthorized. However, when either permissions or passwords are weak, this graphic user interface (GUI) makes it easy for an intruder to find interesting locations he can attack with minimal skill.

- ◆ Every network communication is visible to anyone with a protocol analyzer or sniffer. These tools are software- or hardware-based devices that can capture network traffic and display or record the contents of the communications (packets) sent across the network. This traffic can then be searched for clear-text copies of passwords or other interesting data, including file transfers and email. Sniffers and packet analysis offer network administrators an excellent tool for use in network troubleshooting. Unfortunately, they offer an attacker a rich source of information as well. The captured traffic can also be used to infer situations or intent by noting where the traffic came from and where it's going. For example, seeing a larger than normal amount of traffic from a government to its troop ships might be evidence of some forthcoming activity. (Instructions for movement elsewhere? Attack? Retreat?) To defend against the exposure of confidential information, information should be encrypted. To defend against inference, other methods should be used to disguise the true source of the data, sending fake but plentiful messages at all times to all stations. The use of misinformation can redirect the thought of anyone attempting to use this kind of traffic analysis.
- ◆ The protocols used by the network can have inherent vulnerabilities. Understanding them can give the attacker a way to disrupt communications. TCP/IP, for example, is vulnerable to a number of attacks. Properly designed and configured implementations of this protocol are less likely to become victim to the attacks.

UNDERSTANDING MALICIOUS CODE

Describe and define malicious code.

To understand malicious code, you have to understand its authors, its impact, and the processes that have been developed to deal with it.

Definitions abound in this area, and not all of them are agreed upon. For our purposes, however, we'll state the more common explanations and then explore approaches for dealing with them.

So, Who's a Hacker? What's Malicious Code?

Like any industry, security sports a lot of terms and jargon that everyone is sure they know the meaning of, and yet few can provide good explanations. Following are some explanations for some of them.

Hackers, Crackers, and Phreakers

In the past I've debated the terms *hacker* and *cracker* until blue in the face and I will have to admit defeat. Truly, perception is reality. The term *hacker* will never have its original meaning again. At one time "to hack" meant to attempt to learn how things worked. It was the gleeful exploration of any complicated thing. It could get you into trouble. Like a child who gets burned by playing with fire, hacking systems could crash them, could have unexpected results. Sometimes, however, documentation was sparse, and gurus sparser. The only way to fix something broken, the only way to figure out how to do something, was to just get in there and mess with it. If you were good, you became the guru. You didn't have to hack anymore. You knew. You probably moved on to a system you didn't understand.

Today, however, the term hacking has come to mean malicious exploitation of a system. It means going-to-places-you-aren't-supposed-to-go-to and doing illegitimate things while you are there.

There are some who say *hacking* means experimenting with no malicious intent, while *cracking* means intentional breaking or breaking into, whether for profit or bragging rights. However, the distinction is usually lost in the miasma that is public opinion.

Phreaking, on the other hand has always been the term applied to those who hack into phone systems. This originally started as a way to make free phone calls. Ever more sophisticated devices and software exploits have been developed to hack phones, PBXes (the private branch exchange, or the private phone network within a company, which shares outside lines), and even Telco networks.

NOTE

Looney Tunes? John Draper, also known as Captain Crunch, one of the early phreakers (1971), has now reentered info-security news by starting a security firm. Draper got his nickname when he discovered that the toy whistle included in the Cap'n Crunch cereal boxes could be used to reproduce pay telephone codes and obtain free service. He later developed small electronic "blue boxes" that could be used for the same purpose.

Real Problems and Pseudo Attacks

We often are warned about destructive worms. Malicious code is any code that, either by design or as the result of being run, accomplishes any of the following:

- ◆ Modifies computer programs without the consent of the owner or operator
- ◆ Crashes programs or systems
- ◆ Steals or modifies data
- ◆ Inserts or adds code to a system which might do damage later

On the other hand, we often hear of malicious attacks that turn out not to exist at all. These warnings often seem real and include technical jargon, and a respected sender (someone@someimportanttechfirm.com). Although these hoaxes do not exist and therefore never infect computers, they still can do a lot of damage. They usually request that you send them out to everyone you know in order to prevent these folks from falling victim to the attack. Thousands of unsolicited emails then flood a company's servers and spread to other organizations across the Internet. Such mail storms can clog in-boxes and servers and result in substantial time wastage. The best response to any computer malware warning is to validate its worthiness before passing the information on. This can be done in the following ways:

- ◆ **Checking Internet hoax busting sites**—These sites maintain lists of known hoaxes. One example is www.hoaxbusters.ciac.com.
- ◆ **Checking with well-known alert sites**—such as www.cert.org, www.sans.org, or, if the warning is about a product-specific issue, the security pages on the vendor's Web site. CERT, SANS, and product vendors do not send unsolicited warnings. You can sign up for newsletters and warning lists. If you have any doubt about the veracity of any communication from them you can check the PGP signature. A PGP signature is a digital signature which can be validated by checking against a copy of the user's public key. The vendor's site will have instructions on how to obtain the key to be used for validation.
- ◆ **Reporting the warning to your security department.**

What Protection Does Antivirus Software Provide?

Antivirus software can only protect the computer from known viruses and worms. Why use it then? First, because there are many known viruses and worms still circulating on the Internet, in email, on floppy disks and infected CD-ROMS and yes even in shrink-wrapped product discs. The only way to protect your system from this legacy malware is to obtain, and always run, modern antivirus software on every system in your network. Second, because every reputable antivirus software product offers updates. You see, when a new virus or worm is detected, these companies have a stake in determining the malware's footprint, so they can add that to your current database and enable their product to protect you from the new threat.

In addition, antivirus software for edge servers might offer other services. *Edge servers* are those servers that accept input from untrusted networks and make it available to clients on your trusted network. They also might return responses or requests. Examples are firewalls, mail servers, and Web servers. Antivirus products designed for these servers can block executable attachments from email, filter for malware, and perform other server-specific services. Because many viruses and worms are spread through unsolicited email attachments and attacks on Web servers, it makes sense to have specialized products for these systems. If the attack code never gets to the intended victims, it cannot infect them.

Antivirus products are not failsafe, and they will never completely protect you. But short of removing the floppy and CD-ROM drives, and the network cards, modems, and wireless technologies, there is no other solution.

IMPLEMENTING SYSTEM DEVELOPMENT CONTROLS

Discuss system development controls.

System development controls can be beneficial in two ways: in the use of a strong systems development lifecycle, and in following sound best practices.

System Development Lifecycle

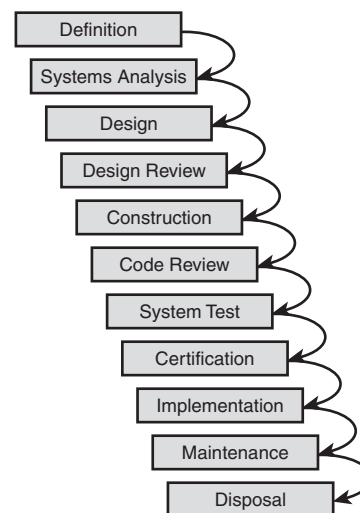
In the beginning, there was chaos. What else would you expect of a new industry? When the first systems were programmed, there was no history of product development to follow. Programmers often followed the “code-and-fix” model of development. The program was written and, if found wanting, was fixed. And fixed again, and then fixed some more. In response to this, the concept of structured software development was devised. Three prominent system development lifecycle models exist. The *waterfall system development lifecycle* model is the best known and has existed for decades. The *spiral systems development lifecycle* is less well known, but might exemplify more closely the model used by organizations with large, albeit younger staff. The process practiced by many newer organizations has been called many things. One name used is RAD, or *Rapid Application Development*.

Waterfall

The classical waterfall approach to software development has been with us for a very long time. Each step from conceptual development to maintenance flows from the top down. Figure 4.8 illustrates the model. Historically, the development process was described as a logical progression of steps. One phase was completed, and then the next phase initiated. Meanwhile, down in the trenches, realists followed the steps, but were not afraid to return to an earlier phase if it meant a better product in the end.

FIGURE 4.8

The waterfall methodology got its name from the way each phase seems to flow into the next.



Chapter 4 APPLICATIONS AND SYSTEMS DEVELOPMENT SECURITY 279

The phases in the process often vary in number as various authors have either expanded the steps or compacted them. The following list is yet another composite:

- ◆ **Conceptual Definition/Feasibility Study**—The need for the software to be developed is described and flushed out during an initial discovery phase. Here, in addition to a needs analysis, a feasibility study can be done. Many projects never proceed past this stage as they are found to be cost prohibitive, or otherwise not a good use of resources. The purpose of the feasibility study is to determine if a business case exists for the system. Several areas are explored, including business (does it solve a business problem?), operations (will it work in our operational model), technical (can it be done?), and financial (do benefits outweigh cost?).
- ◆ **Systems Analysis/Functional Requirements Determination**—Precise descriptions of exactly what is needed. This is done to a fine, granular level of detail. The current system is analyzed to determine what it does, and what should it do, whether through computerization or manual systems. The question is asked, how can it be made better? And changes are recommended on how to solve the business problems associated with this system.
- ◆ **Design/Specifications Development**—A detailed design of how the system will look. It is said that if this is done well, the pseudocode (precise descriptions of the processing with no programming language used) can easily be converted into code with little modification. An example of pseudocode is illustrated in Figure 4.9. If questions arise during this phase that cannot be answered by referral to the functional requirements, the previous phase must be revisited. Two types of design are done: first, a logical design based on user requirements and ignoring any constraints (financial, technical, and so on), and second, a physical design where constraints are considered.
- ◆ **Design Review**—A step-by-step review of the design measuring it against the functional specification. If it is found lacking, a return to the previous phase is necessary.
- ◆ **Construction**—The program is coded according to the design.
- ◆ **Code Review or Walk-through**—Code is reviewed in excruciating detail, step by step to assure the program matches the design.

Print the Customer List

1. Open customer file
2. Do while not end of file
 - a. Read customer number, address, city, state, zip
 - b. Print customer number, address, city, state, zip
 - c. Advance to next record
3. Close customer file

FIGURE 4.9
Pseudocode.

- ◆ **System Test Review**—All aspects of the code are tested looking for functionality, design flaws, and bugs.
- ◆ **Certification/Accreditation**—If the code must meet or is scheduled to meet some formal review for certification or accreditation this is the next step.
- ◆ **Implementation**—Code is put into production. There might be a transition period, where file conversion is accomplished and the old system is changed to the new.
- ◆ **Maintenance**—As errors are found or enhancements required, code is modified, tested, and placed into production.
- ◆ **Disposal**—At some point, legacy code is retired because the system is no longer needed, or has been replaced by completely new systems. For example, a mainframe order entry application might be recoded in and moved to a Web-based interface using PC-based front-end and mid-range systems as the back-end database. The code from the old system is destroyed or archived, but not used in production again.

Although the waterfall system was designed with a phase-to-phase operation in mind, in reality, each phase serves as a control on the others. If a review finds a flaw, or a test a bug, the previous phase can be revisited. In addition, we might best be served if we recognize that each phase does not represent completion for all parts of the program. In other words, in a large development process, some design work might be completing while other work has not begun.

Spiral Lifecycle Model

The waterfall method evokes images of fancy garden waterfalls, or raging mountain streams that cascade down cliff faces and into pools, from which they surge over another drop into yet another pool and on down the mountain. In contrast, the spiral model starts in the middle with the conceptual model of what must be done and spirals outward through its phases, which repeat, at ever widening paths. Figure 4.10 displays this model. The steps it uses are described in Step By Step 4.10.

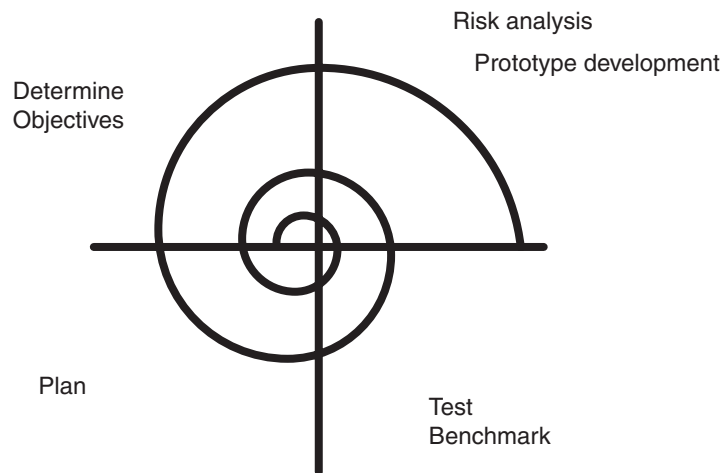


FIGURE 4.10
The spiral lifecycle model.

STEP BY STEP

4.2 Following the Lifecycle Model

1. Develop a preliminary design.

2. Develop a prototype from the design.

3. Develop the next prototype.

4. Evaluate.

5. Define further requirements.

6. Plan and design another prototype.

7. Construct and test this prototype.

8. Repeat steps 3–7 until the customer is satisfied that the prototype meets the requirements.

9. Construct the system.

10. Thoroughly test the final system.

Another spiral model example is presented by the Center for Academic and Research Computing at the University of Missouri: <http://ctr.umkc.edu/~kennethjuwng/spiral.htm>. It describes the spiral model as the waterfall model with the element of risk analysis added. This model is credited to Barry Boehm, chief engineer at TRW, 1988.

In essence, four operations are repeated over and over until the right design is created, which is then put into production. The four operations are

- ◆ **Planning/review**—Determine the objectives of the system to be developed.
- ◆ **Risk analysis, prototype**—First, identify all alternative solutions and perform a risk analysis. Resolve the risks and create the prototype.
- ◆ **Engineering**—Develop and verify the product requirements. Validate the design. Do a detailed design and validate it. Code a test product.
- ◆ **Plan the next phase**—Review for customer satisfaction. Perform requirements planning, development planning, and integration planning, and create a test plan.

Rapid Application Development

Rapid Application Development (RAD) recognizes that the result of software development is a product that meets economic, reliability, and speed-of-development goals. It seeks to develop a product that has 80% of what is desired, but is produced in 20% of the time normally required to meet 100% of the goals. A common saying is that a RAD project has a strong chance of developing the product in the timeframe desired if the company is willing to sacrifice either economy or quality. And, that it has a better chance of achieving its goal if the customer is willing to sacrifice both economy and quality.

The RAD process includes the following stages:

- ◆ High-level end users and designers convene a Joint Application Development meeting. (This is a brainstorming session out of which come the requirements.)
- ◆ Developers build a prototype based on requirements.
- ◆ Designers review the prototype.
- ◆ Customers try out the prototype.
- ◆ A focus-group meeting takes place in which customers and developers refine the requirements and change their requests.
- ◆ A new prototype is developed and the process begins again.

Does this sound something like the spiral development model? The differences here are that the requirements and change review are time-boxed. That is, a limited time is allotted to each phase. As the end of this time approaches, secondary features are dropped to stay on schedule. The repetitive process might function over a day or over a few weeks with the prototype evolving into the operational product. The total time for development might be six months or less. In contrast, the spiral model is not time boxed. It might extend over long periods, and the product is developed with the final prototype as a guide.

The RAD process, if not carefully controlled, can degenerate into quick-and-dirty application development (QADAD). Even its proponents agree that it should not be used to develop an operating system or other product where the need for quality is high, for games where the demand for performance is high, or for any product that is mission- or life-critical.

A more detailed description of RAD can be found at <http://csweb.cs.bgsu.edu/maner/domains/RAD.htm#2>.

Security Control Architecture

A *security control architecture* is the sum of the controls built into the system. It might be controls enforced by the hardware or software. The security architecture for different types of systems will vary. The security architecture of an operating system running on a modern Intel machine can include such features as

- ◆ **Process isolation**—The ability to run different processes and separate them from one another. Each process has its own data and code space; consequently, if a process fails, it can only crash itself, and other running processes are unaffected.
- ◆ **Hardware segmentation**—The isolation of software processes and data via the segmentation of hardware. An example of this can be found in the 80-386 and above Intel systems memory model. In these systems, access to protected and real mode memory address locations is controlled by different memory registers.
- ◆ **Memory protection**—Virtual memory is divided into segments. Each process uses its own segment, and the system keeps its own internal processing separate from that of user mode processing (the running of applications).

Because of segmentation, an unprivileged user process cannot access or modify the memory used by the system.

- ◆ **Least privilege**—Processes have no more privileges than needed to perform functions. For example, only modules that need complete system privilege are located in the kernel (the central area of the operating system), where all essential operations are controlled including memory, disk management, and process management.
- ◆ **Separation of duties**—It is possible to assign privileges on the system so that related privileges are segregated—for example, backup and restore.
- ◆ **Layering**—A structured, hierarchical design of system function. Layers communicate through calls via defined interface.
- ◆ **Security kernel**—Hardware, firmware, and software that implement a reference-monitor concept. A reference monitor mediates access to the system and is protected from modification. Its processing algorithms and implementation can be verified as correct—that is, you can prove that it will always respond as designed.
- ◆ **Modes of operation**—Different system uses are separated into privileged and unprivileged. Access to one does not provide access to the other. Different machine and OS architectures provide different names for this. One name for privileged access, for example, is Supervisor mode; User mode refers to unprivileged access.
- ◆ **Accountability**—With one user per account, you must be able to identify the individual's activity on a system.

A security architecture of a system is the sum of these features. It is important to note that just because a feature is possible, does not mean it is used. The highest security level supported by a system at a particular time is called the *system high* and the lowest, *system low*. Where mandated, a system is tested to ensure that it conforms to the appropriate level for its use. This system of *accreditation* is an official authorization and approval to use the computer system on the network and to process sensitive data. Accreditation, which is a management process, cannot be accomplished until a technical evaluation or *certification* that the system meets security requirements has been done.

You should note that the concept of evaluating security architecture can be extended to networks, and should be evaluated at the network, operating system, database, and file level.

Best Practices

Several systems development practices exert control over the process. These practices can be followed no matter the software development model used.

The first principal concern is the partitioning of development from production. All development work should be done on test systems, not on production systems. Even minor fixes should be done in the development environment, and thoroughly tested before putting the new code into production. This practice minimizes several risks. Because, in some cases, developers must have near-total control over their machines, it is unwise to let these machines be production machines. To allow them administrative control over production machines would be to violate the principle of separation of duties and least privilege. These principles are useful as they avoid potential fraudulent misuse of systems as well as accidental damage, or unauthorized access to sensitive data and processes.

The second promotes documentation of code and of code changes. Good program documentation makes it easier to maintain, and to bring new individuals up to speed faster on the systems. Although false documentation could lead reviewers astray, validated documentation assists reviewers, troubleshooters, and future generations of programmers that must fix or replace code.

The third requires backup of development as well as production code. Many systems are usually in place to back up data and programs to assure business continuity in the face of any disaster. Few have considered the devastating effect of losing source code and code in development.

Fourth, continuous training is essential in a world where rapidly changing and advancing standards, practices, hardware, and methodology means skills can be rapidly outdated.

Finally, the adoption of coding standards, systems development models, practices, and methodologies assists the programming team in producing quality code that is reliable and secure.

USING CODING PRACTICES THAT REDUCE SYSTEM VULNERABILITY

Explain how coding practices and software design can reduce vulnerabilities.

- **Software development methodologies: Are some more secure than others?**
- **Good coding practices prevent flawed software.**

Many argue that it is the programming language that makes a difference in the security of the program. Some claim it is the environment—the combination of operating system, languages, and programming style. Still others say that only “open source” projects can be secure. I’m afraid the argument will never be decided to everyone’s satisfaction. I do believe that everyone will agree: Good software development practices can make more reliable and robust programs with fewer vulnerabilities. In the previous section, we talked about software development methodologies and how adhering to best practices can improve the reliability of software. Now it’s time to talk about application development methodologies, and the effect solid coding practices can have on reducing vulnerabilities.

Software Development Methodologies

Good software can be developed using many different methodologies. Some methodologies can only be performed with certain programming languages. The following major development methodologies are in use today:

- ◆ Structured programming
- ◆ Object-oriented programming
- ◆ Computer-aided software engineering (CASE)
- ◆ Prototyping

Structured Programming

The structured programming methodology was developed in response to the lack of methodology and structure in early development efforts.

Programs were often just massive lists of instructions. Each instruction was executed in sequence until an instruction required a move or jump to another line somewhere else in the code. Execution continued at that point until another branch moved execution elsewhere and so on and so on. This type of programming is very difficult to maintain, makes it difficult to understand what is actually going on, and it is difficult to determine the impact of any changes you might make. It earned for itself the name *spaghetti code*, because of the tangled mess it appears to be. You can still find some of these programs today. I hope you do not have to deal with them. Often these projects had no organization at all. Some of this was due to the early languages and to the lack of training in methodology. Early programmers were often trained on the job, and the emphasis was on syntax, or how to write code that would work, not on making it neat or maintainable.

In contrast, structured programming requires the programmer to be aware of the flow and control of the program.

Structured programming is based on several principles:

- ◆ Modularity
- ◆ Top-down design
- ◆ Limited control structures
- ◆ Limited scope of variables

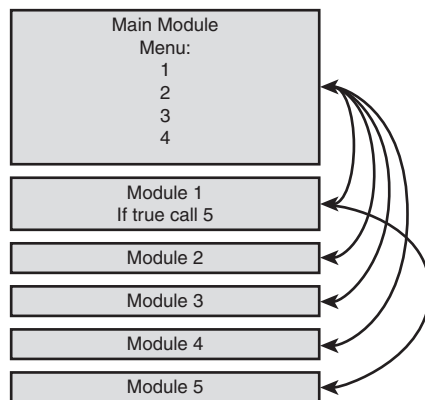
How do you solve large problems? Most people have an easier time solving large problems if they can break them down into small manageable chunks. This is the heart of structured programming.

Instead of composing one large body of code, the work that the program needs to do is broken down into smaller parts that are themselves broken down into still smaller parts and so forth. These parts are called *modules*. Modules are small functional pieces of code that perform a function. Logically you might compare the process to writing a book. You start with a top-level outline which states the topics that will be covered, and then you break each topic down a couple of more levels. The outline can then become the structure within which the words are written that tell the story. Each topic becomes a chapter and its inner levels become subtopics.

Just as the book outline proceeds from a high-level outline to the details, a structured program is based on a *top-down design*. This means a hierarchy of modules branch off a main module.

NOTE

Spaghetti Code Some programmers still use this method of coding today. Many of them are self-taught. It is just as hard to maintain their code today as it was in the past.

**FIGURE 4.11**

Structured programming promotes the use of modules.

The main module is the place where execution of code begins. Each module can also call other modules, but the program eventually returns to the main module either to traverse another path through the program, or to end. If you read a novel, you probably read from one end to the other, but when you use a reference book, you probably look up a topic in the table of contents, and jump to the page or section you want.

Figure 4.11 shows a tiny example of how this might work. You can clearly see the main module and the four choices for program direction. The four modules are also represented; module 1 can also call module 5. The flow of the program might be as illustrated by the arrows which trace the path from the main module to module 1, to module 5, then back to 1 and back to the main module. This is not the only path of execution. It is merely an example of how the activity might flow. In the real world few programs would be this simple; indeed they might have hundreds if not thousands of additional modules.

Although some structured programming languages enable the simple branching statements, structured programming methodology requires more limited control structures. An instruction might require control to go to another statement—that is, the beginning of another module. However, it requires that when that module has completed, control be returned to the calling module. Another example of a control structure used in structured programming is a loop. A loop iterates through a series of instructions and terminates when some condition is met. Perhaps it will continue, adding one to a base number until some present total is matched. Perhaps, it will continue until the user selects the Exit button on the screen. Or perhaps, it continues until the end of a file is found. In the latter case, imagine you are reading a list of names. Each time you read a name, you write it down on another piece of paper. Your loop would look something like Step By Step 4.3.

STEP BY STEP

4.3 An Example of a Simple Loop

1. Read the name.

2. Write it down.

3. Advance to the next page.
 4. Go to step 1: Continue until there are no more names to read.
-

What if a module's purpose is to manipulate some numbers and return an answer? What if the module requires information in order to work? How can data be used within the program? Structured programming insists that data operate within its assigned *scope*. Scope, here, means the part of the program where the data is known and therefore can be utilized. Suppose, for example, you wrote a module to add two numbers. One choice would be to hold the values for the numbers on a global basis, so that the data could be retrieved and manipulated from within every module of the application. We could write a program and make this work, but not every module needs to know the value of these variables. In fact, it is not a good idea to provide global scope, because the data can be manipulated anywhere in the program. How could we ever prove that it was not? However, if we *limited the scope of the variables*, the data could not be modified outside of that scope. To the rest of the modules, the data does not exist.

Object-Oriented Programming

When I drive a car I don't think about the internals of the combustion engine. I don't look under the hood before I open the door and get in. I just don't care. And I suspect many of you don't either. What we want is safe, reliable transportation. (Some of you might be looking for other things but for most of us, it's not the internal workings of the car that matter, but what we and others see on the outside.)

To us then, it is what we can do with the car that matters, not the intimate details of how it works. This is also the essence of object-oriented programming. In an object-oriented program, objects, which are structures that contain data and code, are the building blocks. Just as we make a car take us where we want to go by using the steering wheel, and make it move by pressing an accelerator pedal, objects have an interface by which they are manipulated. Let's look at a simple example.

290 Part I EXAM PREPARATION

Pseudo Code Add Two Numbers

1. declare number1, number2, result as integers
2. main
 - 2.1 number2 = readnumber
 - 2.2 number1 = readnumber
 - 2.3 result = add(number1, number2)
3. add (number1, number2)
 - 3.1 declare sum as integer
 - 3.2 sum = number1 + number2
 - 3.3 return sum

FIGURE 4.12

Pseudo code for adding two numbers.

Class Math

```

Variables:
number1, number2, result, integers
Methods:
Result = add (number1, number2)
  sum integer
  sum = number1 + number2
  return sum

```

FIGURE 4.13

Object-oriented programming: defining the class.

Add two numbers

```

Num1, num2, sum1 = int
get(num1, num2)
Calc = new(math)
Sum1 = calc.add(num1, num2)

```

FIGURE 4.14

Object-oriented programming: adding two numbers.

Our problem, again, is the addition of two numbers. In the structured program, we created a module. Three data variables were used by the module, one for each number and one to return the answer. Inside our module we write the code to do the math. It's easy to trace the execution path by looking at the code we have written and following it along. Figure 4.12 shows simplified pseudocode for such a program.

To use the object-oriented paradigm, we first write a class. A *class* is simply an abstraction, a description of an object. When we actually want to use the code written in the class, we create an object. Figure 4.13 is the pseudocode for our class. As you can see it contains three variables, and its own module, called a method, Add. The code for the Add method simply adds the two numbers it's given and returns the answer. Now, to perform the calculation, we instantiate, or create the object, and then send it a message—or call its method. Figure 4.14 is the pseudocode for this operation.

You might have noticed some similarities here. There are still three variables involved and the code to add them looks the same. There are differences too; the code that actually did the work is hidden from the main program. In structured programming, a module is called to perform a function. In object-oriented programming, an object is sent a message (a command) to perform a function. The function and the data variables are encapsulated within the object. In the structured programming example, the module definition is combined with the instructions which call it. In the object-oriented example, the actual code to add the numbers is further hidden in a separate construct.

The object-oriented concept here is to keep the details hidden. In the real world, programs are much more complex. Objects, like the internal combustion engine of my car, don't need to expose their inner workings in order to be used. We can encapsulate them and only work with them through their exposed interface. For my car, that's a key in the ignition, a steering wheel, and so forth. For the program, it means public methods.

There are other object-oriented concepts as well. Classes, those building blocks of object-oriented programming, can inherit from other classes. *Inheritance*, means that some of the functionality of child classes can come from the parent class. Just as traits such as blue eyes or musical ability are inherited, functionality can be too.

If our previously discussed simple class was made the parent of another class, that other class would inherit the ability to add numbers.

When you first learned addition, you worked with simple integers, 0–10. Later, you worked with larger numbers, and then fractions and decimal numbers. Although the operation was similar, you had to do the addition in a slightly different way. Today, you hardly think of the differences. If you need to add two numbers you simply do so. Your ability to add is *polymorphic*, that is, you can use the word *add* for many purposes. Our simple class, described above, might have been written to accommodate all these types of addition. Its inner workings on how it does this are not relevant to our use of the class. Any correctly structured message add function, whether it includes integers, decimal numbers, or something else, will get us the correct result. (That something else could even be two words. To add two words might mean we obtain their concatenation. So if we enter the two words *red* and *hat*, we obtain the result *redhat*. How this is accomplished is defined by the inner workings, but we don't need to know how the addition is accomplished; we simply need to have a description that tells us this class can be used to add the following types of items. The ability to have one method available for many uses represents the polymorphic characteristic of object-oriented languages.

Computer-Aided Software Engineering

It only seems logical that anything as labor intensive as programming might be automated. *Computer-aided software engineering (CASE)* uses computers to help in the control and management of complex software development projects. CASE tools, programs that have been developed for this purpose, might do anything from keeping repositories of plans, design, code, documentation, and progress to generating actual code. Users of these tools can more quickly get and keep up to speed on project status. They can prevent duplication of effort, possibly translate from design to code and back again, graphically display project progress, and eliminate some of the drudgery of manual project documentation.

Glowing advertisements aside, CASE is not a substitute for a strong methodical approach to software development. Some CASE tools support the structured approach, others are developed to support object-oriented design and programming, whereas still others support visual programming orientations and other methodologies.

NOTE

Want More? For those of you with some programming background, an excellent introduction to object-oriented programming is Peter Mueller's *An Introduction to Object-Oriented Programming Using C++*, which you can find at <http://www.zib.de/Visual/people/mueller/Course/Tutorial/tutorial.html>.

Other good OO (object-oriented) resources are listed at the end of this chapter.

CASE methodology often emphasizes customer involvement promoting the use of focus groups and prototypes.

Impacting Security Through Good Software Design and Coding Practices

Anyone can write a computer program. That's true, the basic concepts are not hard, and tools exist which enable linking together of already created components with a few pieces of sample code. Writing a program that does what it's supposed to do, and that is reliable and secure, is a much harder thing to do. The software methodologies discussed previously, and many others, seek to alleviate the problem that occurs when no methodology is used.

In addition to following some methodology, there are distinct coding practices that, if followed, can virtually eliminate many of the reliability and security issues that we have today. Why aren't they used? There are many reasons:

- ◆ The software market has been driven by the twin philosophies of “first to market” and “feature rich.” This means that development time is spent on features, not security and reliability, and that products are rushed, and therefore testing time is minimized.
- ◆ Modern operating systems are developed to allow a large, diverse, number of devices. This means that it must accommodate a large range of software device drivers which are written by other companies.
- ◆ As consumers we expect a large amount of diverse software that will run on our systems, and we also demand backward compatibility. Many problems are not the fault of the operating system but of the software it must accommodate.
- ◆ The attitude of the many software developers and many software development companies (and supported by the market) has been that quick development of software that does only part of what has been promised is okay. The companies feel that the additional promised features can be made to work via patches, service packs, and the next release.

Chapter 4 APPLICATIONS AND SYSTEMS DEVELOPMENT SECURITY 293

- ◆ In the rush to be first to market and flush with features, software development has abandoned some of the rules, standards, and techniques developed over the years.
- ◆ As consumers, we have come to accept these software development flaws as the norm.
- ◆ The complex nature of software means it is difficult to eliminate errors and vulnerabilities in a reasonable timeframe and to therefore produce affordable products.
- ◆ We now have a heightened sensitivity and awareness of the problems that vulnerabilities in software can create due to the massive connectivity and exposure our systems now have. More connectivity means more exposure to danger. Not too many years ago, few systems were publicly exposed to the Internet, and few people had Internet access, thus many software flaws were never taken advantage of.
- ◆ An ethically poor attitude and the availability of prewritten attack code. Anyone can find prescribed code and even GUI-based programs that can be used to attack systems.

It is not my purpose to lump all software developers and all software development companies and all products and projects into this mold; instead I mean to emphasize that no one single problem has led to the larger number of attacks and identification of vulnerabilities exposed in software. There are many problems; well-written software will not solve them all. However, there are techniques that lead to software with fewer vulnerabilities and greater reliability. Although many of these techniques have been known for a long time, two recent books have documented them. These books are *Building Secure Software*, by John Viega and Gary McGraw, and *Writing Secure Code*, written by Michael Howard and David LeBlanc. Here are some of the techniques they detail:

- ◆ **Eliminating buffer overflows**—Buffer overflows represent more than 50% of the security advisories. Buffer overruns exist when data must be entered or passed to modules. The module does not check the incoming data to see whether it fits within the “buffer” or area of memory set aside for the data. For example, a buffer overrun, or overflow, can occur if space has been reserved to accept the two-character abbreviation for U.S. state names and someone enters the entire state name.

index	numbers
0	4
1	7
2	10
3	3
4	8
5	22
6	25
7	23
8	67
9	45

FIGURE 4.15

An array is a structure that holds data in order. The data can be referenced in a software program by indicating the position it fills in the array.

If the code does not check to ensure that only two characters are entered, the result is a buffer overflow. Some buffer overflows are relatively harmless; they merely crash the program (obviously I jest). Others can give an attacker the opportunity to execute further attack code, eventually giving them root access on the system. Buffer overflows can be eliminated by coding practices which test data entry and by programs that search code for potential buffer-overflow problems.

- ◆ **Prevent array indexing errors**—An *array* is an ordered data structure used in programming to hold several pieces of data. It can be an array of characters, numbers, or other types of data. You can think of an array like the mailboxes at a post office. Each box has a number and mail is sorted into the boxes according to these numbers. Because the boxes are numbered and arranged in order, it is easy to locate any box. Likewise, array elements can be located and data stored or retrieved from any position by referencing its number or index. Figure 4.15 illustrates an array of numbers. To print the number 22, the programmer would reference the array name and the index 5. Software errors occur when the programmer makes a mistake in referencing elements of the array. Different programming languages number the elements in the array differently; some start the index at 0 and others at 1. Thus, for an array of five elements, the last element might have an index of 4 or 5 and the first element might have an index of 0 or of 1. Improper referencing can cause the program to “fall off” the end of the array and produce unpredictable results. Proper coding techniques prevent errors because these types of errors are tested for, and *bounds checking*, or making sure there are no references to nonexistent array members, is done within the program.
- ◆ **Utilizing good access control**—Access-control techniques are available to the programmer. The operating systems for which they code offer granularity in protecting files, printers, and other types of data. When the programmer ignores or abuses these capabilities, he does not allow the administrator to enforce them. In both Unix and Windows NT and above, file access controls can be set in the file system. They can be set administratively either through a GUI or through commands, but they can also be set programmatically. The overall design of the project should specify the minimal access necessary for code and user and the programmer should follow these specifications.

Chapter 4 APPLICATIONS AND SYSTEMS DEVELOPMENT SECURITY 295

- ◆ **Principle of least privilege**—It is much easier to write code which runs in a security context that allows any privilege and access anywhere. Instead, however, code should be written that operates at the least privilege level required to do the job. Web servers that run on Unix systems are usually written so they can execute at a normal user level. However, Microsoft IIS server runs in the context of the operating system. Which type of Web server do you think operates to meet this security design principle?
- ◆ **Defense in depth**—Looking for the holy grail of software security is always tempting. What design, coding technique, language, or methodology will produce the securest code? Unfortunately, there is no silver bullet. Just as we lock valuables in a safe even though we can lock our hotel room, so must we apply several security principles. Secure coding practices, designs that operate with least privilege, proper access controls, and many other techniques need to be applied in order to ensure the best security.
- ◆ **Hiding secrets**—To prevent unauthorized access to data and systems, authentication and authorization techniques are used that rely on the user possessing some secret, such as a password. Other times, encryption is used to secure data. Although these techniques make systems and data more secure, they introduce new problems. To validate the user's knowledge of the password, the password must be stored on the system, or some other technique must be used. To enable encryption and decryption, the key or keys must be stored. It's difficult to store these types of secrets on systems. Sooner or later, an attacker will use them. Numerous techniques have been developed, such as storing one-way encrypted hashes of passwords, further encrypting the keys, obfuscating the keys, storing them on hardware separate from the computer, and so on. The issue is that secrets are hard to hide, but there are techniques and approaches, that can either make them harder to access or make them useless if accessed with anything other than the right credentials. These techniques can vary depending on the languages and operating systems used.
- ◆ **Remember the weakest link**—Strong cryptography can help secure systems. But attackers will always look for the weakest link.

What good, for example, are steel barred doors, if windows are easy to open, and glass to break? Why attack the password file hoping to discover the administrator's password when a buffer overrun exploit can gain the attacker control over the system? Time spent looking for and securing the weak links is well spent. Attackers many times can be rebuffed when the known weak links are secured.

Good design and coding practices can mean better, more reliable and more secure software. The results are quantifiable. Where they are implemented, the number of bugs is reduced and customer satisfaction improves.

CASE STUDY: TRUSTWORTHY COMPUTING

ESSENCE OF THE CASE

The essence of this case and the thrust of Trustworthy computing is

- ▶ **Availability**—Lack of system outages, and self-recovery when necessary.
- ▶ **Security**—Data and systems should be protected.
- ▶ **Privacy**—Users control their own data.
- ▶ **Trustworthiness**—From chips to customer service, a broad category that means customers can rely on Microsoft products.
- ▶ **Manageability**—Relative to size and complexity, the system is easy to install and maintain.
- ▶ **Responsiveness**—The company takes responsibility for its product and helps customers to resolve problems.
- ▶ **Transparency**—The company is open with customers.
- ▶ **Accuracy**—Results are free from error, and protected.
- ▶ **Usability**—Software is easy to use.

SCENARIO

Can software development processes be changed to provide more secure code? You are all currently involved in just such a project. In January 2002, an internal memo was leaked to the press. It outlined an internal project that sought to produce more secure code.

Described as a necessary change of attitude to assist the company in producing “trustworthy computing,” the memo from Chief Software Architect, Bill Gates, asked all employees to participate in the project. You can read what is purported to be the January memo at <http://www.computerbytesman.com/security/billsmemo.htm>.

CASE STUDY: TRUSTWORTHY COMPUTING

ANALYSIS

To its credit, Microsoft has not indicated that this is an easy task that can be solved by a couple of months of code review and programmer training. Further explanation of the long-term (10–15 year) commitment necessary for the success of the vision, and the necessity that all organizations participate, is illustrated in a later whitepaper delivered by Craig Mundie, Senior Vice President and CTO, Advanced Strategies and Policy. You can read this paper at <http://www.microsoft.com/presspass/exec/craig/05-01trustworthywp.asp>.

Many were quick to criticize the memo as just a marketing ploy. Microsoft has been heavily criticized for a long time for producing security-weak, buggy products. This memo was seen as an attempt to change public attitude without doing anything. Microsoft also announced an immediate month-long shut down of work on .NET, the next version of the Windows operating system. The announced purpose was the training of programmers on writing secure code and the scouring of .NET and other existing product code for software bugs. Various sources at Microsoft claim some 9,000 programmers have been trained and that the shutdown lasted for two months.

Additionally numerous bugs have been corrected and the orientation of .NET changed to focus on security versus features.

In response to the original memo and later announcements, a Web site, www.trustworthycomputing.com, put up a page to refer to a www.google.com search page for “Microsoft security or privacy flaw or flaws or hole or holes.” News of this Web page initially dominated the press response to Microsoft’s campaign.

In contrast, vendors who have promoted “trusted systems” engineered to deliver security solutions are seizing the opportunity to advertise their solutions. On-board smart card readers in keyboards, and other hardware devices, as well as specialized BIOS-level routines are touted as the answer in the April 4, 2002, article “Signs of Trustworthy Computing,” available at <http://www.wired.com/news/business/0,1367,51521,00.html>.

Trustworthy Computing is a goal that might not be accomplished for many years, if ever. However, there cannot help but be improvements in computer security along the way.

CHAPTER SUMMARY**KEY TERMS**

- Basic input output system (BIOS)
- Blended malware
- Boot sector virus
- Brute-force attack
- Cache
- Centralized controlled computing
- Centralized systems
- Data consistency
- Data independence
- Data mining
- Data recovery
- Data redundancy
- Data reuse
- Data warehouse
- Decentralized
- Dictionary attack
- Distributed
- DMBS
- Dynamic random access memory (DRAM)
- Flooding net
- Grid computing
- Hardware segmentation
- Hierarchical database

Applications can contribute to the security of our computer systems or continue to add additional vulnerabilities to them. The choice is ours. We must scrutinize the applications that will be used on our systems and within our networks, and we must not forget the application development process and its contribution to security or vulnerability. In addition, we should realize the impact of the Internet, or chats, channels, and email as portals for the distribution of malicious applications as well as harmless ones. It is no longer enough to manage the applications that are part of our organizations' business processes. We must realize how easy it is for peripheral code to enter our systems for good or evil.

CHAPTER SUMMARY

- Knowledge-based systems
- Logic bomb
- Macro virus
- Malware
- Multi-partite virus
- Object-oriented
- Persistence
- Polymorphic virus
- Primary key
- Primary storage
- Process isolation
- Random access memory
- Rapid application development
- Real memory
- Referential data integrity
- Registers
- Relational database
- SANs
- Secondary storage
- Security control architecture
- Security controls
- Semantic data integrity
- Sequential access
- Spoofing
- Static random access memory
- System development lifecycle
- Time of Check to Time of Use (TOC/TOU)
- Trap door
- Trojan horse
- Virtual memory
- Virus
- Web services
- Worm

APPLY YOUR KNOWLEDGE**Exercises****4.1 Password Cracking**

How easy is it to crack passwords? It's certainly easy to talk about the reasons for strong passwords, and the need to develop alternatives to them. But how much of a problem is it really? To find out, obtain and run a password cracker on a system on which you are authorized to do so. The easiest process to follow is to set up a test Windows NT or Windows 2000 system, create accounts and populate them with passwords, and then run a cracking program against them. This exercise details how to do so.

Estimated Time: 1 hour

1. Locate a system capable of running Windows 2000. *Do not utilize a production system!* Not only is it unethical to crack passwords on a system, it is illegal. You could find yourself in serious trouble. Cracking passwords as part of an audit to determine the use of strong passwords is a legitimate security technique; however, when doing so, permission must be obtained in writing. For our purposes, it is only necessary to demonstrate the technique, not to perform a true audit.
2. Load Windows 2000 Professional or Server. If you do not have a licensed copy for testing purposes, you can usually obtain a limited use (time-bombed) demonstration copy. This system will only be used for this experiment and therefore only needs to be operational for a few days.
3. Download a 15-day trial copy of LC4 from <http://www.atstake.com/research/lc/download.html>. This is the latest version of the popular and notorious Lophcrack product from @stake.
4. Populate the Windows 2000 system with at least a dozen user accounts. This is done by selecting Start, Programs, Administrative Tools, Computer Management, Local Users and Groups, Users and selecting New users.
5. Create passwords for the users which reflect typical choices by users—for example, names, birthdates, popular characters, pet names, and so on, as well as some strong passwords (those including upper- and lowercase letters, numerals, and punctuation marks).
6. Install LC4. (You must be logged on with an administrative account.) To install only requires double-clicking on the downloaded executable and accepting the defaults.
7. Run LCR. Select Start, Programs, LC4, LC4.
8. At the LC4 wizard welcome page, click Next.
9. On the Get Encrypted Passwords page, leave the default option, Retrieve from Local Machine, checked and click Next.
10. On the Choose Auditing Methods page, leave the default option, Strong Password Audit, checked and click Next.
11. On the Pick Reporting Style page, leave the defaults alone and click Next. Click Finish.
12. Let the password cracker run for some time. Note the passwords cracked and the time it takes to crack them.
13. To end the program, from the File menu, select Exit.

You should read the help files for LC4 and understand that the brute-force capability of the trial copy is not functional. Strong passwords that would eventually be cracked using the brute force techniques will be not be cracked using the trial program.

APPLY YOUR KNOWLEDGE

After your experiment consider: How could you use a password cracking program in a security program?

Review Questions

1. Give an example of a distributed software environment.
 2. Give an example of a non-distributed software environment.
 3. Why do distributed systems increase the risk quotient of software systems?
 4. Explain the difference between worms, virus, and logic bombs.
 5. Discuss the difference between a relational database and an object-oriented database.
 6. Why are distributed database systems harder to protect?
 7. How can a paging file pose a risk to systems?
 8. Does a SANS pose any special security risk?
 9. Name and define two types of software attacks.
 10. Why is remote administrative software dangerous?
2. A protocol analyzer is an example of a what?
 - A. Virus
 - B. Hacker tool
 - C. Legitimate network administration tool
 - D. Trojan horse
 3. Which of the following is *not* a legitimate way to deal with an announcement of circulating malicious code?
 - A. Check with CERT.org.
 - B. Check with your security officer.
 - C. Do a search on the Internet for hoax busting sites.
 - D. Forward the notice to all of your friends.
 4. Which of the following is true about antivirus programs?
 - A. They facilitate secure remote administration.
 - B. They can be configured to block executable attachments from email.
 - C. They discover and destroy or quarantine all virus attacks on computers on which they are installed.
 - D. They rebuff attacks from malicious code.

Exam Questions

1. Back Orifice is an example of a what?
 - A. Remote administration tool
 - B. Logic bomb
 - C. Virus
 - D. Trojan horse
5. A software development methodology which uses extensive prototyping and is best suited for applications where economy or quality might be sacrificed is which of the following?
 - A. Spiral
 - B. Waterfall
 - C. Unstructured
 - D. RAD

APPLY YOUR KNOWLEDGE

6. The waterfall methodology of software development is characterized by which of the following?
 - A. A progression of steps. Each step must be completed before the next one can follow.
 - B. An iterative pattern in which planning and risk analysis dominate.
 - C. Characterized by focus groups, prototyping and time-boxing.
 - D. A loose application of methodology in which programmer style is more important than documentation or formal practices.
7. The ability to work in a higher level view of a problem is called what?
 - A. Abstraction
 - B. Layering
 - C. Data hiding
 - D. System high
8. A software development methodology characterized by modularity, data hiding, and limited control structures is called which of the following?
 - A. Object-oriented programming
 - B. Structured programming
 - C. Computer-aided software engineering
 - D. Spaghetti code
9. The problem with using cryptography to hide information is which of the following?
 - A. Every cryptographic system can be broken.
 - B. It's difficult to implement.
 - C. Hiding the keys that decrypt the data is a more difficult thing to do.

- D. Few inexpensive programs exist that enable this to be done on today's PC systems.

Answers to Review Questions

1. An e-commerce site with a database back end. For more information, see the "Distributed Systems" section.
2. A standalone database accessed by terminals. For more information, see the "Non-Distributed Systems" section.
3. One way in which distributed systems increase the risk quotient of software systems is that they offer more opportunities for the spread of malware. Viruses can be spread by removable storage on any system, but distributed systems can also be infected by email, access to Web sites, chat rooms, and use of instant messenger programs. See the section "Malware for Distributed Systems" for more information.
4. Worms spread themselves by traveling from computer to computer. Viruses hide their code within other, legitimate programs. A Trojan horse is a malware program that disguises itself as something else. See the section "Malware for Distributed Systems" for more information.
5. A relational database stores its data in tables composed of rows and columns. The tables are "related" by relationships between the primary key of one table and the foreign key of another. The SQL language is used to query the database to store and retrieve data. An object-oriented database typically stores its data by mapping its objects into tables. The object-oriented programming language is used to store and retrieve data. See the "Data Models" section for more information.

APPLY YOUR KNOWLEDGE

6. Distributed databases are harder to protect because the data can be itself distributed across multiple locations. Transactions can involve access to and manipulation of data in more than one database and therefore, there are problems with consistency. See the “Database Issues” section for more information.
7. A paging file is used to temporarily store data to disk during processing. Data is paged in and out of memory to disk, thus extending the memory space available. Unfortunately, if sensitive data, such as unencrypted data or plaintext passwords, exist in memory, they can also be paged to disk. Although the paging file is protected when the system is running, when a system is shut down it is not. After shutdown, the paging file exists on disk as an ordinary file. If the paging file was not cleared at shutdown, the sensitive data exists on disk. Booting the system to another OS might expose the sensitive data. See the section “Storage and Storage Systems” for more information.
8. A SANs can pose a security risk because security is often not designed in. Although operating systems can have access-control designed in, a SANs that is accessible from all systems cannot have any special controls available. This might have been less of an issue when SANs systems were contained in the data center and used lesser-known communications channels, but SANs systems are now becoming distributed systems and migrating to IP. See the section “Storage Area Networks.”
9. Software attacks can be dictionary attacks, brute-force attacks, spoofing, man-in-the-middle, sniffing, scanning, and so on. See the section “Attacking Software.”
10. Remote administration software can be used to administer systems from locations other than the protected data center. If an unauthorized person can obtain a legitimate account with privileges to run them, they can attack from a remote location. See the section “Illegitimate Use of Legitimate Software.”

Answers to Exam Questions

1. **D.** Back Orifice is a Trojan horse. This software was developed to remotely control systems without permission. The “server” portion of the product is often innocently installed by an administrator who has been tricked into doing so. The “client” is installed on the system used to attack the victim. Also, many admins have been tricked into thinking this is a legitimate product and installed the system thinking to use it for their own work, only to find a backdoor has allowed an unauthorized individual to control their systems. See the section “Illegitimate Use of Legitimate Software.”
2. **C.** A protocol analyzer is an example of a legitimate network administration tool. It is used to troubleshoot networking problems. It can be used by an attacker to inspect traffic on the network. See the section “Network Software.”
3. **D.** Forwarding the notice to all of your friends only perpetuates the hoax, if that is what it is. By checking with official resources (CERT, your security department) you might discover the true nature of the problem and how to deal with it (ignore, patch). If the nature of the threat is unknown, contacting your security department will ensure its investigation and proper action.

APPLY YOUR KNOWLEDGE

Getting others excited about a nonexistent problem is in itself a problem as it clutters up the mail servers and can reduce the availability of network resources. See the section “Real Problems and Pseudo Attacks.”

4. **B.** Can be configured to block executable attachments from email. This feature is present in most antivirus programs that are made for email servers. By eliminating executable attachments, a rich source of malware is prevented from reaching the end user. Because it is difficult to train users not to click attachments, preventing attachments from reaching users eliminates a threat. See the section “What Protection Does Antivirus Software Provide?”
5. **D.** Rapid Application Development is a methodology that seeks to bring projects to fruition quickly. It is difficult to do so without sacrificing something. See the section “Rapid Application Development.”
6. **A.** A progression of steps. Each one “flows” down to the next, hence the name. See the section “System Development Lifecycle.”
7. **A.** Abstraction is the ability to view a problem from a high, conceptual level. See the section “Security Control Architecture.”
8. **B.** Structure programming. See the section “Structured Programming.”
9. **C.** Hiding the keys is problematic. Although writing cryptographic code is difficult, many software development environments include prewritten interfaces that simplify its use. Although it is true that, eventually, encryption might be broken, an attacker will first seek to obtain the keys. (Why do the difficult thing, when the easy solution exists?) See the section “Impacting Security Through Good Software Design and Coding Practices.”

APPLY YOUR KNOWLEDGE**Suggested Readings and Resources**

1. Anderson, Ross. *Security Engineering*. Wiley, 2001.
2. Grimes, Roger A. *Malicious Mobile Code*. O'Reilly, 2001.
3. Howard, Michael, and David LeBlanc. *Writing Secure Code*. Microsoft Press, 2001.
4. Krehnke M.E., and D. K. Bradley. "Data Marts and Data Warehouses: Keys to the Future or Keys to the Kingdom." In *Handbook of Information Security Management, Fourth Edition*, edited by Micki Krause and Harold F. Tipton. Auerbach, 2001.
5. McConnell, Steve. *Code Complete*. Microsoft Press, 1993.
6. Vallabhaneni, S. Rao. *CISSP Examination Textbooks*. SRV Professional Publications, 2000.
7. Viega, John, and Gary McGraw. *Building Secure Software*. Addison-Wesley, 2002.
8. Whitehead, Katherine. *Component-Based Development*. Addison Wesley, 2002.
9. <http://catalog.com/softinfo/objects.html> ("What Is Object Oriented Software?" by Terry Montlick).
10. <http://msdn.microsoft.com/vstudio/techinfo/documentation/default.asp> (Microsoft Visual Studio).
11. <http://www.atstake.com/research/lc/download.html>.
12. <http://www.CERT.org>.
13. <http://www.codagen.com> (Gen-it Architect, Codagen code generation).
14. <http://www.computer.privateweb.at/judith/> (A medical expert system).
15. <http://www.computerbytesman.com/security/billsmemo.htm>.
16. <http://www.cyberdyne-object-sys.com/oofaq2/> (many basic and detailed explanations).
17. <http://www.emsl.pnl.gov:2080/proj/neuron/kbs/demos.html> (Knowledge-based systems).
18. <http://www.emsl.pnl.gov:2080/proj/neuron/neural/what.html> ("What Is an Artificial Neural Network?").
19. <http://www.entercept.com>.
20. <http://www.globus.org/research/papers/anatomy.pdf> (Grid computing).
21. <http://www.hoaxbusters.ciac.com>.
22. http://www.lwfug.org/~abartoli/virus-writing-HOWTO/_html/ (Linux virus writing how-to).
23. <http://www.malware.org>.
24. <http://www.methods-tools.com/> (Software Methods and Tools—a source for information on software development methods and tools).
25. <http://www.microsoft.com/presspass/exec/craig/05-01trustworthywp.asp>.
26. <http://www.okena.com>.

continues

APPLY YOUR KNOWLEDGE**Suggested Readings and Resources** *continued*

- | | |
|---|---|
| 27. http://www.omg.org/gettingstarted/gettingstartedindex.htm . | 34. http://www.telelogic.com/about/apart.cfm (Visual modeling tools Telelogic software test and project management tools). |
| 28. http://www.onesecure.com . | 35. http://www.trustworthycomputing.org . |
| 29. http://www.rational.com (Rational Rose, Rational Software). | 36. http://www.wired.com/news/business/0,1367,51521,00.html . |
| 30. http://www.research.ibm.com/massive/ (IBM Massively distributed systems). | 37. http://www.wired.com/wired/5.11/heartof.html ("Heart of Darkness," a report on Bulgarian virus writers). |
| 31. http://www.sans.org . | 38. http://www3.ibm.com/software/ad/vajava/ (Visual age for Java). |
| 32. http://www.sdmagazine.com/documents/s=4077/sdm0203f/0203f.htm . | |

OBJECTIVES

Discuss the uses of cryptography including confidentiality, integrity, authentication, and nonrepudiation.

- ▶ Cryptography is not an easy subject to study. Many complicated mathematical algorithms exist, and few people who are not dedicated to the field find pleasure in examining them. The first step in approaching a study of cryptography is to understand what it is used for and some of the common terms. After you understand how cryptography is used and how critical it is to computer security, you'll be ready to face those complicated algorithms.

Compare and contrast symmetric and asymmetric algorithms.

- ▶ Two major types of cryptographic algorithms—symmetric and asymmetric—exist. Understanding how they work and their weaknesses and strengths is critical to understanding how to use them to protect, not expose, sensitive data and resources.

Describe PKI and key management.

- ▶ PKI is the cornerstone of much that is new in the use of encryption technology today. It is also being touted as the new hope—PKI can solve all our computer security issues. This is, of course, not true. Even though PKI presents tremendous opportunities for securing data, if it's improperly implemented and used, it is just another good thing gone bad.

Detail common methods of attacking encryption, including general and specific attacks.

- ▶ If you do not know how encryption is usually attacked, you can fall victim to the theory that some forms of encryption are not hackable. This theory is very wrong. If you understand common methods of attack, you can assist in designing strong networks that are resistant to these types of attacks.




CHAPTER 5

Cryptography



OUTLINE

Introduction	310	PKI and Key Management	318
Uses of Cryptography	310	Methods of Attack	319
Confidentiality	310	General Attacks	320
Integrity	311	Ciphertext-Only Attack	321
Authentication	311	Known-Plaintext Attacks	321
Nonrepudiation	312	Chosen-Plaintext Attacks	322
		Chosen-Ciphertext Attacks	322
		Specific Attacks	322
		Brute-Force	322
		Replay Attacks	323
		Man-in-the-Middle Attacks	323
		Meet-in-the-Middle Attacks	324
		Birthday	325
Cryptographic Concepts, Methodologies, and Practices	313	Chapter Summary	327
Symmetric Algorithms	313		
Asymmetric Algorithms	315	Apply Your Knowledge	329
Message Authentication	316		
Hash Functions	316		
Digital Signatures	317		
Key Length	317		
One-Time Ciphers	318		



STUDY STRATEGIES

- ▶ Read the introductory information to get a high level understanding of the key components.
- ▶ Read the entire chapter concentrating in on the key technical areas.
- ▶ Go through the chapter concentrating on the exercises and understanding how all of the pieces fit together.



“The Cryptography domain addresses the principles, means, and methods of disguising information to ensure its integrity, confidentiality, and authenticity.

The candidate will be expected to know basic concepts within cryptography; public and private key algorithms in terms of their applications and uses; algorithm construction, key distribution and management, and methods of attack; and the applications, construction, and use of digital signatures to provide authenticity of electronic transactions and nonrepudiation of the parties involved.”

—Common Body of Knowledge study guide

INTRODUCTION

There is no silver bullet when it comes to network security. One technology comes close, however: cryptography. Most people do not understand how cryptography works and why it is important that it become a critical part of their security arsenal. This chapter introduces the key concepts that are needed to be able to use and integrate security into your environment.

USES OF CRYPTOGRAPHY

Discuss the uses of cryptography including confidentiality, integrity, authentication, and nonrepudiation.

Cryptography (abbreviated *crypto*) can be used for a variety of purposes to protect information. When most people think of crypto, they think of making sure no one else can read a certain piece of information; keeping their secrets secret. This plays a key role in crypto, but there are actually four other main goals of cryptography. Each of these is discussed in the following sections.

Confidentiality

Confidentiality is preventing, detecting, or deterring unauthorized access to information. I have sensitive data and I want no one else to be able to read it. This is a fundamental goal of encryption.

What is important to remember is that not all encryption provides confidentiality. Some encryption schemes provide only integrity and authentication without providing confidentiality. The reason this is a key point is most people instantly associate encryption with confidentiality and that can be a dangerous assumption to make under certain circumstances. Confidentiality of information can be obtained through both symmetric and asymmetric encryption.

Integrity

Integrity is preventing, verifying, and detecting the alteration of data or information you have sent. You have to make sure that someone cannot modify your information without your knowledge. Some people ask why this is a separate category, because they would argue that you cannot modify information if you cannot read it. If the information is protected from a confidentiality attack and unreadable how could someone modify the information? The answer is, “Very simply.” You just need to find out the value of a field that you know and use that as a starting point to modify information you might not know.

Let’s look at an example to make this clearer. If an employee gains access to the spreadsheet that human resources maintains to keep salary information, the employee cannot read the salary information because that field is encrypted. However, the other fields are not encrypted so if the employee knows that the CIO of the company makes more money than he does, he could copy the encrypted value for the CIO’s salary and paste it in his own field. This employee might not know the value to which he changed his salary, but as long as it was higher than his initial salary, he would consider the attack a success. This is one example where you can modify information even if you cannot read it. Hash algorithms are typically used to provide for integrity of information.

Authentication

Authentication involves identifying an individual or verifying that the individual is part of a certain group. For example if you try to get into a bar, the bouncer does not really care who you are as a person;

he just wants to make sure you belong to that group of people who are 21 or older. In other cases if you are trying to use a credit card, the merchant wants to make sure that you are the person who is listed on the front of the card. You typically can authenticate someone based on one of three attributes:

- ◆ Something the person knows, such as a password
- ◆ Something the person has, such as a token
- ◆ Something the person is, or biometrics

Encryption is used by all three authentication methods. No matter what you use to authenticate, you want to make sure the information is protected as it travels the network and that it is also secure when it resides on the backend server. If an attacker can intercept a password when it crosses the network or on the backend server, she could impersonate that user on the system. This is also possible with biometrics because after a biometric reader assesses your physical attributes, that information is sent and stored in binary format. If someone could intercept the binary format, she could impersonate that user on the system.

Nonrepudiation

Nonrepudiation is critical when it comes to digital signatures. It deals with proving in a court of law that someone was the originator. E-commerce would never have taken off if a merchant could not prove that someone was the originator of the transaction. In traditional contracts our signature serves as proof that we contractually obligated ourselves to an agreement. Because that signature is unique to you, someone at a later point in time can prove that you committed yourself to that agreement, meaning you cannot *repudiate* it, or get out of it.

This same type of proof needs to be obtained in the digital world. Otherwise, people could place orders and if, a day or two later, the price decreased, they could deny that they ever placed the order. If this could occur, no one would use the Internet for any type of e-commerce. Nonrepudiation is a feature of asymmetric encryption that allows you to prove that someone actually sent a message. It is equivalent to an actual signature.

CRYPTOGRAPHIC CONCEPTS, METHODOLOGIES, AND PRACTICES

Compare and contrast symmetric and asymmetric algorithms.

As previously discussed, cryptography has several properties and no single technique can achieve them all. By putting various different pieces together, you can achieve a strong robust solution. When talking about cryptography, the following basic terms need to be defined:

- ◆ **Plain text**—A message in its original form. Remember that any type of message can be encrypted. So even though the word has *text* in its name, plaintext is really a generic term and can refer to an executable, a zipped file, a word-processor document, a spreadsheet, or any type of information you would want to keep protected and secure. This is the data before anything has been done to it.
- ◆ **Ciphertext**—A message after it has been encrypted.
- ◆ **Encryption**—The process of taking a plaintext message and converting it to ciphertext.
- ◆ **Decryption**—The process of taking ciphertext and converting it back to a plaintext message. The key thing with encryption and decryption is this: If you take a plaintext message, convert it to ciphertext, and then decrypt it back to plaintext, the plaintext, decrypted message must match the original plaintext message that was input into the encryption algorithm.

Symmetric Algorithms

Symmetric encryption is often called *single-key* or *secret-key encryption*. That is because a single key is used for both encryption and decryption of the information. So if I wanted to send you an encrypted message using symmetric key encryption, I would encrypt the message with a key, and send you the key and the message. You would then use the same key to decrypt the message. The key thing to remember is that the key has to be kept secret. Whoever knows the key not only can decrypt messages but also can encrypt messages to impersonate the sender. As you can tell from the previous sentence, the logistics create a problem.

314 Part I EXAM PREPARATION

If I am sending you an encrypted message, it means that the media I am using to transport the message is not secure. If it was secure, I would not need to encrypt the message. But if the media is not secure and we do not have a secure link, how am I going to get you the key? This is one problem with symmetric-key encryption; the key must be sent over a secure channel. If someone can intercept the key, they can read the information.

The other problem with symmetric key encryption is nonrepudiation. If we are both using the same key, how can one of us prove in a court of law that the other one sent the message? Let's look at an example. Alice wants to send a secure message to Bob using symmetric encryption. She sends the encrypted message and then sends him the key over a secure channel. Bob decrypts the message. Two weeks later, Alice denies ever sending the message, so Bob tries to take legal action against Alice. Alice claims that she never sent the message. Her argument is that because Bob had the same key she has, Bob wrote a message looking like it came from Alice and encrypted it. Because they both have the key, Bob has no way of proving she actually sent the message.

DES (data encryption standard) and triple DES are the most popular symmetric key encryption schemes used. Because DES uses a 56-bit key, based on current computer speeds it is no longer considered secure; a brute-force attack can be performed in a short period of time. Triple-DES uses a larger key length and is the symmetric algorithm of choice. However, things are going to change because AES (advanced encryption standard) is being developed by the National Institute of Standards and Technology (NIST) to replace DES (visit www.nist.gov). A new algorithm was selected via a NIST-sponsored contest. The algorithm that won is Rijndael. Remember that there is no way to prove an algorithm is secure except by letting a bunch of really smart people beat on it for a long period of time. Even though Rijndael is still being tested, NIST has announced its selection. FIPS-197 is the official newly approved government standard that defines the Rijndael algorithm (<http://csrc.nist.gov/encryption/aes/> and <http://csrc.nist.gov/encryption/aes/frn-fips197.pdf>). Initial feedback indicates that it is a solid algorithm and will become the next big standard for symmetric encryption.

Asymmetric Algorithms

Asymmetric encryption is often called *two-key encryption* or *public-key encryption*. It involves two keys: a public and a private key. The public key is given to anyone who wants it and the private key is kept secret by the user. Anything that is encrypted with one key can only be decrypted with the other key. To make sure that no one can read your message to Bob, you would encrypt the message with Bob's public key. Bob would then use his private key to decrypt the message. Anyone along the path would be unable to read the message. Even if they were able to intercept Bob's public key they still could not read the message. Remember that after a message is encrypted with Bob's public key, the public key cannot decrypt it. The only way to decrypt it is by using Bob's private key, which only he should have. So with asymmetric encryption the public key does not have to be sent over a secure channel but it must be sent over a trusted channel. Otherwise an attacker could generate a fake key for Bob and send it to you.

One of the drawbacks of symmetric encryption was that it did not address nonrepudiation. Asymmetric handles nonrepudiation very eloquently. Remember the sentence earlier about asymmetric encryption; anything that is encrypted with one of the keys can only be decrypted by the other. What happens if I encrypt a message with my private key? It can only be decrypted with my public key. So if Alice encrypts a message with her private key, anyone can read the message because anyone has access to her public key, so it does not address confidentiality. However, when Bob receives the message and successfully decrypts it with Alice's public key he has determined that the only person that could have created this message is the person that has Alice's private key; because Alice is the only one who has access to her key, we just proved that she sent the message.

You might be thinking that it is great you can get confidentiality if you encrypt with someone's public key and you can get nonrepudiation if you encrypt with my private key, but how do you get both confidentiality and nonrepudiation? Easy, you perform two steps. First, you would encrypt a message to Bob with your private key and then you would encrypt the output with Bob's public key. Now what is sent across the wire is secure. Bob would decrypt with his private key to read the message and then decrypt with the your public key to prove that you sent the message.

If asymmetric encryption is so powerful, why do you need symmetric encryption? The reason is speed. Symmetric encryption is very fast and asymmetric encryption is very slow. So in practice, for confidentiality, most messages are encrypted with symmetric encryption and use asymmetric encryption as the secure channel. Alice and Bob have successfully exchanged public keys and they want to send a secret message. Alice could just encrypt the entire message with Bob's public key but because this message is very large, this would be inefficient because the algorithm is very slow. Instead Alice would generate a secret key and use that to encrypt the message. She would then take the secret key, which is very small, encrypt it with Bob's public key, attach it to the message, and send both pieces together. Bob would decrypt the key portion with his private key, obtain the key, and use it to decrypt the rest of the message.

RSA is the asymmetric algorithm of choice and is used in most implementations that utilize this type of encryption.

Message Authentication

Message authentication codes (MACs) are used to make sure the message has not changed in transit and therefore protect it against integrity attacks. Authentication codes can be very basic or complex but they perform some checks to determine whether any of the information has been modified. A basic check that is not secure is parity checks. Parity checks the number of 1's in the message before it was sent and the receiver checks the number of 1's when it is received to make sure they match. So if a single bit is modified this will catch it but if two bits are modified it will not.

The basic operation is that a check is performed on the message before it is sent and attached to the message. The receiver will perform the same calculation and check the results to make sure they match. If they match, the message is processed; if they do not match, the message is dropped and an error is generated.

Hash Functions

A *hash function* is a one-way transformation that cannot be reversed. It takes input data and produces a smaller fixed length output.

Having the output, there is no way to figure out what the original input text is. Another characteristic of strong encryption is there should be no way to pick two input data streams that produce the same output. Hash functions are very popular with digital signatures because they reduce the amount of information that has to be encrypted. The most common implementation of hash functions is MD5.

Digital Signatures

Digital signatures are used to ensure nonrepudiation. Previously, when discussing asymmetric encryption, we discussed how encrypting with someone's private key can ensure nonrepudiation. However remember that asymmetric encryption is very slow, so encrypting the entire message would be very inefficient. Instead, the message is first put into a hash function. A hash function takes a message of any length and produces a smaller fixed length output. So by using the hash function, we decreased the size of the message. This smaller message is then encrypted with the private key of the sender.

Key Length

A common rule of encryption is that all encryption is breakable; it is just a matter of time before it's broken. It might take 200 years, but by utilizing a brute-force attack, which is an attack that tries every possible key, the encryption will eventually be broken. The amount of time it takes to perform a brute-force attack depends on key length. The longer the key, the more possible potential values for the key, which means it will take longer to guess. For example, if we are talking binary numbers, a key length of two can be broken very quickly because there are only four possible combinations. (2 to the power of 2 equals 4.) However, jumping to a key length of 56 bits gives 72,057,594,037,927,936 possible keys. This is derived by raising the number 2 to the power of 56, 2^{56} . Because computers are binary devices a 56-bit key is composed of 56 bits and each bit can either be zero or one. So you can quickly see the longer the key length the longer it will take to break the encryption.

The rule of thumb is that the usefulness of the information should be less than the time it takes to brute force the encryption. For example if one company is going to buy another company within three months, the first company wants to keep this information private.

After the first company buys the second, however, this information will become public and no longer needs to be protected. So the usefulness of this information is three months. If the company uses a key length that can be broken within 12 months, that works fine for this information. However, if the information is about a new airplane that can go to the moon, and it will take 20 years to build this airplane, a much stronger encryption must be utilized in order to keep the information safe from the public.

Another important point is that computers are constantly increasing in power and speed. Just because it takes 10 years today to break a certain type of encryption does not mean a year from now it will not take less than a year. Thus, you are really shooting at a moving target when you deal with key lengths.

One-Time Ciphers

A *one-time cipher* is often considered to be unbreakable encryption. That is not really a completely accurate statement. The reason people make this claim is each time you encrypt a message you use a new key. So you would never ever use the same key twice. Now even if someone was able to perform a brute-force attack and break the encryption, it would only let them read that one message and no other message. So it is a very strong form of encryption, but it requires the user to maintain a list of keys so it can use a different one each time. In reality for one-time ciphers the user carries around a hardware device that generates a new key every minute.

PKI AND KEY MANAGEMENT

Describe PKI and key management.

As we start talking about encryption, one of the key principles is that the secrecy of encryption is based on the secrecy of the key, not the secrecy of the algorithm. When using asymmetric or symmetric encryption, you need to have keys in order to encrypt or decrypt the information. To communicate with a couple of people, managing keys yourself is easy, but what happens when you roll out encryption across a large enterprise. Requiring everyone to manage their own keys would get out of hand very quickly.

So in these cases you create a centralized authority for managing keys. This central server is called a *public key infrastructure* server or PKI server and is used to manage public keys of various individuals and companies. However, it needs to store more than just keys. When we talked about asymmetric encryption, we mentioned that the keys do not need to be sent over a secure channel but they need to be sent over a trusted channel. You have to make sure though that when someone says, “Hi, I am Bob and this is my key,” the person is who he says he is. The way you achieve this trust with PKI is through *digital certificates*, which we refer to in this chapter as certificates. There are certificate authorities who sign and issue certificates validating that you are who you say you are. When a person or a company obtains a certificate, they have to show physical proof that they are the entity they are claiming to be. After they prove this, the certificate authority will sign the certificate. Several certificate authorities such as Verisign perform this function across the Internet. When a company sets up its PKI server, the company would obtain the public key for a certificate authority through trusted means. Now when someone presents a key and certificate to the PKI, it can validate the signature of the authority and verify that it is legitimate.

What happens if a certificate needs to be revoked? When this occurs, the certificate authority maintains a list of certificates that have expired, been revoked, or are no longer valid for one reason or another. This list is maintained but is not pushed out to PKI’s because it would consume too much bandwidth and not be efficient, so instead a pull model is used. Periodically, it is up to the PKI to pull down the latest list from the certificate authority so that it can determine whether a certificate it receives is valid. The name of this list is the certificate revocation list.

METHODS OF ATTACK

Detail common methods of attacking encryption including general and specific attacks.

As discussed, there are various encryption techniques that can be used to protect your information. But how do you know that the encryption techniques are robust and really doing what they say they are doing? How do you know that there are not hidden backdoors in the program that someone can use to extract information?

The simple answer is that we do not know how robust a given technique is when it is initially developed. When it comes to encryption, there is no mathematical proof that can be performed that will tell you an encryption scheme is secure. The only way to know the strength of an encryption scheme is to let the world examine it and then attempt to break its cipher. This would normally be performed over an extended period of time before the code is accepted as a secure means of communication across an unsecured network.

That is why a new technique that has only been around for a couple of years is considered, untested, and therefore not secure. With encryption, something is considered unsecure until it has been proven that it cannot be broken by a bunch of really smart people. These people whose goal is to crack encryption are called *cryptanalysts*. Only after cryptanalysts have unsuccessfully tried to break a scheme for three to five years, do people consider the encryption scheme secure.

In this section we look at various ways to attack encryption schemes. The first group consists of general attacks that can be performed against encryption. The second group involves specific attacks that people use to break encryption. In most cases breaking encryption involves finding the key that was used to encrypt the data. After you know the key, you can decrypt the data and read the encrypted message. With encryption, the secrecy of the encrypted text is based on the secrecy of the key, not the secrecy of the algorithm. This means that even if someone knows the algorithm, without the key they cannot crack the encrypted text. Therefore it is fairly common for the algorithm to be open and published because if it is done correctly, it will not make it any easier to crack the encrypted message.

General Attacks

Four general attacks can be performed against encrypted information:

- ◆ Ciphertext only
- ◆ Known plaintext
- ◆ Chosen plaintext
- ◆ Chosen ciphertext

As you move down the list, the attacks become easier to perform. This should not be surprising because as you move down the list you are given more information on which to base your analysis. The more information you are given to solve a problem, the easier it becomes. We will look at all of these in detail but in most cases you are only given the ciphertext. The other attacks are more appropriate if you also compromise someone's machine or in a lab environment.

Ciphertext-Only Attack

With a ciphertext-only attack (COA), the only thing the cryptanalyst has is encrypted text. This is your traditional attack because if you are using encryption to protect your data over a non-secure link, it is assumed that someone will be able to intercept the encrypted text. The whole purpose of encryption is if someone obtains your encrypted text, they cannot read your original message. So this type of attack is very difficult with strong encryption algorithms. *Strong encryption* refers to algorithms that have stood the test of time and no one has found a means to defeat it.

A critical point to cover is that all encryption is breakable, it is just a matter of time. Brute-force attacks are always possible. This is where you try every possible combination until you find the proper key. A critical point with brute-force attacks is, how do you know when you successfully cracked the key? With binary data gibberish, the actual data could look very similar to the encrypted information. Brute-force attacks are discussed in the "Special Attacks" section later in this chapter.

Known-Plaintext Attacks

Known-plaintext (KPA) attacks imply that for a given message the cryptanalyst somehow was able to find the original plaintext message that was used to generate the ciphertext. Two parties might be using the same key and algorithms for several messages and the goal is to find the key. For one particular message the cryptanalyst now has the plaintext message and the corresponding ciphertext. This attack depends on whether there are patterns between the two and the overall strength of the algorithm. Finding plaintext for a given message could make it much easier to crack the key or keep the difficulty level the same. Also the overall length of the message would dictate how valuable or successful this attack will be.

For example, let's imagine that we are using a basic substitution algorithm. Each letter in the alphabet is substituted for another letter. There is a one-to-one mapping. Now a known-plaintext attack would tell you the mapping for every letter that appears in the message. If the message is short, it might only reveal 20% of the key, but if the message is long it might reveal 90% of the key. After you have that much of the key, it is easy to obtain the rest of the key.

Chosen-Plaintext Attacks

In some cases, access to the device that generates the encryption can be obtained without obtaining the key. In this case, you could feed in whatever plaintext you want and receive the corresponding ciphertext. This is one step easier than the known plaintext. With that attack, a cryptanalyst could not pick the plaintext; they are at the mercy of the system. With this attack, they can now pick whatever plaintext they want. The chosen plaintext would contain every single letter in the alphabet. By doing this, the attacker would obtain the mapping for every character and therefore you obtain the key.

Chosen-Ciphertext Attacks

The last general attack is a very sophisticated attack. In this attack, you can pick the ciphertext and the system will give you the corresponding plaintext. As you can imagine, by doing this you can obtain a lot of critical information that would make it easier to crack a given algorithm. However this attack is considered theoretical, and in most cases is only possible in a lab. In normal operations the chances of performing such an attack are very slim, probably nil.

Specific Attacks

In this section we will look at specific attacks that can be launched against encryption systems.

Brute-Force

As we mentioned earlier, all encryption is crackable, it is just a matter of time. So if a vendor tells you that it has proprietary encryption that is uncrackable, run for the hills because the vendor is lying to you.

First, the strength of encryption is based on the *secrecy of the key* not the secrecy of the algorithm. So the only reason you would keep an algorithm proprietary is if it wasn't any good. Second, remember all encryption can be cracked from a brute force standpoint. Because the goal is to find the key you could go and try every possible combination. If the key was composed of letters you would try every possible combination. The beginning of such an attack would look like A, AA, AB, AC, and so on. Eventually you will find the key. It could take 500 years to find it, but it could still be cracked. Therefore when we pick a key length, we have to figure out the time it would take to brute-force that key length and make sure the information content expires before the technique can be brute-forced. For example, if I only have to keep something secret for two days, encryption that could be cracked by a brute-force attack in two weeks would work fine. However, if the value of information has to be kept secret for 10 years, two weeks would be too short a period of time.

Replay Attacks

A *replay attack* involves taking encrypted information and playing it back at a later point in time. For example, to gain access to a network a user would enter a password which is sent over the wire encrypted to the server. You cannot read the password because it has been encrypted with a large key. However, you would sniff the encrypted password and when you want to impersonate a given user, you would just reply or send the server the encrypted information you gathered off the network. The best way to defeat replay attacks is to put some piece of information like time into the equation. So if you try to replay information 10 minutes from now it would not work because the time factor would not match for the data you are trying to replay.

Man-in-the-Middle Attacks

When we talked about symmetric and asymmetric encryption, we said that symmetric keys have to be sent over a secure channel but asymmetric keys only have to be sent over a trusted channel, not necessarily a secure channel. The reason a trusted channel is needed is to prevent an attacker from inserting themselves in the middle of a communication channel and impersonating both sides. For example, say that Alice and Bob want to communicate using asymmetric encryption.

NOTE

Crack Crack is a program written to crack the encryption that is used to store passwords on Unix operating systems. It was originally written to crack the crypt encryption which is a variant of DES used to encrypt Unix passwords. Essentially, crypt used the password as the key and encrypted a set string to produce the ciphertext. Then, when someone entered her password, it would decrypt with the password the user entered and if it returned the set string the user knew the password was valid; if it did not then the system denied access. Crack is pretty basic compared to today's cracking programs, but when it first came out it was very powerful and it showed the impact that all encryption is crackable; it is just a matter of time.

They exchange keys, but they do so on a non-trusted communication media. Evil Eve controls the router or an access point that all of the traffic flows through, so she has inserted herself in the middle of the communication. Now Eve would generate a false public-private key pair for both Alice and Bob. Now when Alice and Bob try to exchange keys, she intercepts the real keys and sends the fake keys to Alice and Bob respectively. Now because Eve controls the keys she can decrypt, modify, and re-encrypt all information that is sent between the two parties. Alice and Bob think they have valid keys because they did not bother to send them through a trusted source or channel.

Meet-in-the-Middle Attacks

Most people have heard of DES and 3-DES or triple DES, but have you ever heard of double DES? What is wrong with double DES that caused the developers to go right to triple DES instead? The reason has to do with a potential vulnerability that exists with double DES; the attack is called a meet-in-the-middle attack. Essentially when you do the first round of encryption, you encrypt the message with key1 to yield ciphertext 1, which is shown in the following formula:

$$E(M, K1) = C1$$

Then you encrypt ciphertext 1 with key 2 to yield ciphertext 2, which is shown in the following formula:

$$E(C1, K2) = C2$$

Now to perform a meet-in-the-middle attack, you need to have both the plaintext message and the ciphertext, so you can already see that this is not a practical attack in most situations. The way this works is you try all possible keys to try and yield C1 with $E(M, K1) = C1$; then you start from the other end and try to decrypt C2 with all possible keys to yield C1 with $D(C2, K2) = C1$. Now all possible keys for k1 are 2^{56} because DES uses a 56-bit key. All possible keys for k2 are 2^{56} also. Now if you add $2^{56} + 2^{56}$, it yields 2^{57} which means because of this weakness double DES only gives an effective key length of 57 bits which is only one more than DES. So, because of this, cryptographers skipped double-DES and went straight to triple-DES instead.

Birthday

When dealing with hash functions, because they are a one-way function, it is critical that the chances of two random messages hashing to the same value is slim. It should also be difficult if not impossible to figure out that the input text was based solely on the output text. The birthday attack against hash functions deals with trying to find two different messages that hash to the same value. If this can be found, information could be implied and potential weaknesses could be found. The name derives from the birthday game which involves taking a room full of people and figuring out the chance that two people have the same birthday. Originally you would think because there are 365 possible birthdays that with a small group of fewer than 100 people the chance of two people having the same birthday would be extremely low; in reality, though, the number is quite high, actually greater than 50%. So the lesson to be learned is even though there are a high number of possible values that something or someone can take on, the chances of two having the same value are extremely high even if the range of answers contains a lot of values.

CASE STUDY: ENCRYPTION CAN BE A DOUBLE-EDGED SWORD

ESSENCE OF THE CASE

This case is an interesting one. The essence of the case involves the following:

- ▶ **A strong system for file encryption**—A very necessary part of business is being able to keep information confidential. Having a file encryption process available to all would seem to be a boon.
- ▶ **Human nature is human nature**—Why read the documentation? Why understand what you are doing?

SCENARIO

Windows 2000 and Windows XP include a free file encryption utility—the Encrypting File System (EFS). This tool is built in. To use it, an encryption/decryption bit is set on a file or a folder. If the bit is set on a folder, all files placed in the folder are encrypted. A combination of random symmetric keys (used to encrypt the file) and symmetric keys (used to protect the encryption key) are used. The symmetric keys are bound to the user account via a self-signed certificate (unless certificate services are established). The user's public key is used to protect the file encryption key, and the user's private key can be used to decrypt the encryption key, which then is used to decrypt the file.

continues

CASE STUDY: ENCRYPTION CAN BE A DOUBLE-EDGED SWORD*continued*

- ▶ **Solutions exist**—Such as disabling EFS until a PKI can be established to ensure the availability of recovery agents.
- ▶ **Systems in a domain might not be vulnerable because a domain-level recovery agent is available**—However, historically, things have happened to corrupt or remove this key as well.
- ▶ **The most vulnerable users to this issue are the very ones who will use it and be caught**—This includes the home user, the small business person, and the company without central data systems with domains and experienced technical people.

In Windows 2000, a file recovery agent exists and can also decrypt the file. Window XP systems not in a domain do not have a file recovery agent.

Unfortunately, users of EFS often receive no training and few if any read the documentation that clearly states the user's keys must be archived to provide backup should the original keys become corrupt. The keys are stored in the user's profile (a collection of configuration information and folders that reside by default on the user's hard drive). Should anything happen to the profile, the keys can be lost or damaged.

Users do not generally archive their keys, and it is not practical for company system/network administrators to do so for them (there is no automated way to do so and thousands of users would mean thousands of archived keys and no key management system). This means that something as simple as a corrupt profile, disk error, or disk crash can destroy the keys. When a user's machine is fixed (drive replaced, profile regenerated, system reinstalled, and so on), even if the encrypted files are backed up or still present, they cannot be decrypted because keys are missing. With luck, the recovery agent can be used to recover the files; however, many people have lost access to critical, sensitive files due to this problem.

CASE STUDY: ENCRYPTION CAN BE A DOUBLE-EDGED SWORD

ANALYSIS

Why would such a product exist? Why don't people read directions? Shouldn't our data processing gurus know about these things? Here is a case where solid encryption has risen on its hind legs and bitten the ones who use it. Worse, few are taking the steps to properly manage it. If steps are taken to archive keys, ensure recovery agents exist, and train users and administrators, this is a good system to use.

Likewise, if an organization decides it does not want the trouble, it is easy to disable this tool to prevent the user from encrypting files with EFS. Instead, it's turned on by default, and it's easy enough to implement. Therefore, a real threat of danger exists. In fact, one consultant I know receives at least one new case a week where someone has encrypted a file and then the keys have been destroyed and their data lost.

CHAPTER SUMMARY

Cryptography plays a key role in obtaining security for an organization. It does not solve all of the world's problems but plays a key role in defense in-depth across an organization. Especially now that organizations are connected to untrusted networks like the Internet, it is critical that people take measures to protect their information. e-Commerce dictates that you must be able to protect information, validate the accuracy of information, and prove that an entity actually sent a message. All of these goals need cryptography to be achieved. Having a good understanding of the different algorithms and the pros and cons of each is critical for any security professional.

KEY TERMS

- Advanced encryption standard (AES)
- Asymmetric encryption algorithm
- Authentication
- Birthday attack
- Brute force attack
- Certificate authority
- Chosen ciphertext (CCA)
- Chosen plaintext attack (CPA)
- Ciphertext
- Ciphertext only attack (COA)
- Confidentiality
- Cryptanalyst
- Cryptography

CHAPTER SUMMARY

- Data encryption standard (DES)
- Decryption
- Digital certificate
- Digital signature
- Encryption
- FIPS-197
- Hash function
- Integrity
- Known plaintext attack (KPA)
- Man-in-the middle attack
- Meet-in-the-middle attack
- Message authentication code (MAC)
- Nonrepudiation
- One-time cipher
- Plaintext
- Private key
- Public key
- Public key infrastructure (PKI)
- Replay attack
- Rijndael
- Symmetric encryption algorithm
- Triple-DES (3DES)

APPLY YOUR KNOWLEDGE**Exercises****5.1 Disabling EFS on a Windows 2000 Professional Computer**

If EFS is not used in your environment, it should be disabled. This is easy to do. The following instructions are for a Windows 2000 Professional computer.

Estimated Time: 5 minutes

1. Open Start\Programs\Administrative Tools\Local Security Policy.
2. Navigate to and expand the Public Key Policies container.
3. Select the Encrypted Data Recovery Agents container.
4. Right-click the certificate in the details pane labeled `file recovery` and select Delete. In Windows 2000, when no file recovery agent exists, file encryption cannot take place. (This is not true in Windows XP. Windows XP Professional requires a different process to disable EFS.)
5. Right-click the Encrypted Data Recovery Agents container and select Delete Policy. This prevents the inclusion of another recover certificate at a later date without the creation of a new policy.
6. Close the Local Security Policy.

Review Questions

1. Discuss the difference between confidentiality, integrity, and authentication.
2. How is a digital signature useful in an e-commerce transaction?

3. Explain the difference between symmetric and asymmetric encryption algorithms.
4. List and explain two problems with symmetric algorithms.
5. A message encrypted with the public key belonging to Jane and sent to her over the network is captured by Peter. Because the public key is publicly available, what prevents Peter from decrypting and reading the message meant for Jane?
6. Asymmetric algorithms can be used to produce nonrepudiation. How is this accomplished? Why is it true?
7. Why isn't public key encryption used for all encryption purposes?
8. Why is it that we say a longer key provides better protection from being broken?
9. What does a cryptanalyst do? Why?

Exam Questions

1. The message in its original form is an example of what?
 - A. Plaintext
 - B. Ciphertext
 - C. Cleartext
 - D. Hash
2. Which of the following is NOT an example of a symmetric key encryption algorithm?
 - A. Rijndael
 - B. DES
 - C. 3DES
 - D. RSA

APPLY YOUR KNOWLEDGE

3. Bob wants to send a private message to Mary and wants no one else to be able to read it. He also wants Mary to be able to know that it came from him. He both signs and seals (encrypts) the message. The following keys are used in which manner?
 - A. Bob uses Mary's public key to encrypt the message and his own private key to sign it.
 - B. Bob uses Mary's private key to encrypt the message and his own public key to sign it.
 - C. Bob uses Mary's public key to encrypt the message and his own public key to sign it.
 - D. Bob uses Mary's private key to encrypt the message and her public key to sign it.
4. A one-way transformation that cannot be reversed is a what?
 - A. MAC
 - B. Hash
 - C. Ciphertext
 - D. Plaintext
5. A way to establish that a key belongs to a particular user is to use which of the following?
 - A. One-time cipher
 - B. Digital certificate
 - C. Digital signature
 - D. Hash
6. A type of cryptographic attack in which the device that generates the encryption is obtained but not the key is a what?
 - A. Chosen-ciphertext attack
 - B. Plaintext attack

- C. Ciphertext only attack
 - D. Chosen-plaintext attack
7. Which of the following is a type of attack in which encrypted information is taken and played back at a later point in time?
 - A. Replay attack
 - B. Brute-force attack
 - C. Man-in-the-middle attack
 - D. Meet-in-the-middle attack

Answers to Review Questions

1. Confidentiality is the prevention, detection, or deterring of unauthorized access to information. Authentication is proving that you are who you say you are, and integrity is preventing, verifying, and detecting the alteration of data. See the sections "Confidentiality," "Integrity," and "Authentication" for more information.
2. The digital signature serves as proof that a specific individual participated in a transaction. The purchaser cannot deny that he has ordered the item. This feature of digital signatures is non-repudiation. See the section "Nonrepudiation" for more information.
3. Symmetric encryption algorithms use a single key, which can both encrypt and decrypt the plaintext. Asymmetric encryption algorithms, on the other hand, use a matched pair of keys. If one key is used to encrypt, the other one must be used to decrypt. See the sections "Symmetric Algorithms" and "Asymmetric Algorithms" for more information.

APPLY YOUR KNOWLEDGE

4. One problem is that the use of a single key creates the problem of key distribution. I must somehow get to you the key I used to encrypt the message. In addition, if I want to share multiple messages with multiple people, we each need to share multiple keys. Another problem is that a single key cannot be used for nonrepudiation. Because the key is shared, its use cannot prove that a specific person used it. See the section “Symmetric Algorithms” for more information.
5. When data is encrypted with the public key of a public/private key pair, only the private key can be used to decrypt it. The public key will not work. Because the private key is kept by Jane, only Jane, when she receives the message, will be able to decrypt it. See the section “Asymmetric Algorithms” for more information.
6. Asymmetric algorithms use two keys. To digitally sign something, Jane’s private key is used. When the message is received, the signature can be proven to belong to Jane because only Jane’s public key can decrypt it. Furthermore, because only Jane has her private key, only Jane could have signed the message; therefore nonrepudiation exists—Jane cannot deny that she signed the message. See the section “Asymmetric Algorithms” for more information.
7. Public key encryption is very slow, so most uses of it use private key encryption to encrypt the cleartext and use public key encryption to encrypt the private key that must be sent to the recipient. See the section “Asymmetric Algorithms” for more information.
8. All encryption is breakable; the object is to make it take a long time. Because all data in the computer is binary, a small key presents only a few possible combinations of 0s and 1s. A larger key presents a lot more. If a brute-force algorithm, which tries every possible combination, is used then it is logical that a larger key, with more possible combinations, will take longer to crack. See the section “Key length” for more information.
9. A cryptanalyst attempts to crack encryption algorithms. A new encryption algorithm must be tested (by trying to crack it) for many years before it can be presumed to be secure. Cryptanalysts do this work. See the section “Methods of Attack” for more information.

Answers to Exam Questions




1. **A.** Answer B is the encrypted plaintext. Answer C is a font style. Answer D is also wrong. See the section “Cryptographic Concepts, Methodologies, and Practices” for more information.
2. **D.** Answers B and C are incorrect because they are symmetric key encryption standards of the U.S. government. Answer A is the new U.S. standard, so it’s also incorrect. See the section “Symmetric Algorithms” for more information.
3. **A.** Answers B and D are wrong because Bob does not have access to Mary’s private key. Answer C is wrong because Mary cannot use his private key to decrypt the signature. See the section “Asymmetric Algorithms” for more information.

APPLY YOUR KNOWLEDGE

4. **B.** Answer A is a message authentication code or check used to determine whether a message has been changed in transit. Answer C is incorrect because it's the encrypted plaintext. Answer D is incorrect because it is the message before it is encrypted. See the sections "Message Authentication" and "Hash Functions" for more information.
5. **B.** The digital certificate binds the key to the user entity. Answer A is a type of encryption algorithm that must use a new key each time, so it's incorrect. Answer C is a digital signature used to determine who sent the message, so it's incorrect. Answer D is a type of one-way encryption algorithm, so it's incorrect. See the sections "One-Time Ciphers," "Hash Functions," "Asymmetric Algorithms," and "PKI and Key Management" for more information.
6. **D.** Answer A, chosen-ciphertext attack, is one where you pick a ciphertext and get a corresponding plaintext. Answer B is an attack in which you know the original message. Answer C is one in which you only have the ciphertext. See the section "General Attacks" for more information.
7. **A.** Answer B is an attack in which every possible combination is tried, so it's incorrect. Answer C is where an attacker inserts himself into the middle of a communication channel and impersonates both sides, so it's incorrect. Answer D is a special attack based on the vulnerability of double-DES, so it's incorrect. See the section "Specific Attacks" for more information.

APPLY YOUR KNOWLEDGE**Suggested Readings and Resources**

1. Atreya, Hammond, Paine, Starrett, and Wu. *Digital Signatures*. RSA Press, McGraw Hill, 2002.
2. Frankel, Sheila. *Demystifying the IPSec Puzzle*. Artech House, 2001.
3. Ganapathi, S.J. "Fingerprint Authentication: Shifting the Electronic Security Paradigm." www.scmagazine.com, February, 2002.
4. Gove, Ronald A. "Fundamentals of Cryptography and Encryption." In *Handbook of Information Security Management*, edited by Micki Krause and Harold Tipton, Auerbach, 1999.
5. Heiser, Jay. "Introduction to Encryption." In *Handbook of Information Security Management, Fourth Edition*, Volume 2, edited by Micki Krause and Harold Tipton, Auerbach, 2001.
6. Kahn, David. *The Code Breakers: The Story of Secret Writing*. Scribner, 1996.
7. Murray, William Hugh. "Principles and Applications of Key Management." In *Handbook of Information Security Management*, edited by Micki Krause and Harold Tipton, Auerbach, 1999.
8. Schneier, Bruce. *Applied Cryptography, Protocols, Algorithms and Source Code in C, Second Edition*. John Wiley and Sons, 1995.
9. Schneier, Bruce. *Secrets and Lies, Digital Security in a Networked World*. Wiley, 2000.
10. Vallabhaneni, S. Rao. Chapter 5, "Cryptography." In *CISSP Examination Textbooks*, Volume 1. SRV Publications, 2000.
11. <http://www.cryptography.com/> (home of Cryptography Research, Inc. It has links to conference papers, articles on protocols, and crypto author sites).
12. <http://www.faqs.org/faqs/cryptography-faq/> (cryptography FAQ that includes a series of articles which define cryptography topics).



OBJECTIVES

Explain the difference between public versus government requirements for security architecture and models.

- ▶ Understanding the differences in requirements between governments and public entities will aid in your understanding of the security models and architectures that exist.

Discuss examples of security models including the following:

- **Bell-LaPadula**
 - **Biba**
 - **Clark-Wilson**
 - **Access Control Lists**
- ▶ These security models conceptually define how access to resources on systems may be controlled. They also offer opportunities for understanding systems that you may have no experience with. The more that you know of different models, the better you will be able to choose the right model for a current architecture choice, or for one that was in existence before your involvement.

Explain the basics of security architecture.

- ▶ Understanding the security architecture of a system is important for understanding how to secure it. Learning basic concepts and terms is a start whether your intention is to participate in a formal evaluation, select evaluated products, or merely to understand the systems with which you work. To secure systems, it is first necessary to know what security functionality they have. To determine functionalities, you have to study the security architecture. Using a recognized security architecture evaluated product may save some time. Understanding that evaluation, and what you must do to meet it, will allow you to have more secure products in place.



CHAPTER 6

Security Architecture and Models

OBJECTIVES

Describe and contrast information system security standards including:

- **Trusted Computer System Evaluation Criteria (TCSEC)**
- **Information Technology Security Evaluation Criteria (ITSEC)**
- **Common Criteria**

- ▶ Although Common Criteria is the recognized security standard today, many products exist that were evaluated by previous standards (TCSEC and ITSEC for example). Therefore, it is important to know something about these standards as well. In addition, even if it is not in your power to specify or purchase evaluated products, understanding the criteria that are considered to make systems secure will allow you to better understand and secure the products that you do have.

Describe the Internet Protocol Security (IPSec) standard.

- ▶ The TCP/IP protocol has no security built in. IPSec provides that. It is an Internet Engineering Task Force (IETF) standard, and yet multiple products exist with varying interpretations of the standard. Although Domain 3 addresses the technical aspects of networking, it is important here to view the standard's architecture and how it can be used.

OUTLINE

Introduction	338
Requirements for Security Architecture and Models	340
Security Models	342
Bell-LaPadula	342
Biba	345
Clark-Wilson Model	346
Access Control Lists	347
A Review of the Security Models	347
Security System Architecture	348
Reference Monitor	348
Open Versus Closed Systems	350
Security Principles	351
Security Modes	352
Labels Versus Access Control Lists	353
Covert Channel	354
Information System Security Standards	355
TCSEC—The Orange Book and the Rainbow Series	356
Orange Book Classifications	357
Criticisms of Orange Book	358
Rainbow Series	359

OUTLINE

Information Technology Security Evaluation Criteria	360
Differences Between the Orange Book and ITSEC	361
The United Kingdom Information Technology Security Evaluation and Certification Scheme	361
Common Criteria	362
What Is Common Criteria?	363
Part 1: Introduction and General Model	364
Part 2: Security Functional Requirements	365
Part 3: Security Assurance Requirements	367
Evaluation Assurance Packages or Levels	368
Areas Not Addressed by the Common Criteria	369
A Comparison of the Orange Book, ITSEC, and Common Criteria	370
IPSec	370
Uses for IPSec	371
Architectural Components of IPSec	372
Chapter Summary	375
Apply Your Knowledge	377

STUDY STRATEGIES

- ▶ The best way to study security architecture is to use this chapter to obtain an overview of the topic and then apply the steps later in the chapter to make the topic become more than a dry listing of criteria.
- ▶ Get a copy of one of the three standards for detailed study. Each of them offers comprehensive information on what makes a system secure. You will find a review of security policy, features, components, and assurance. The objective here is not to know in intimate detail what each level of each standard requires. The objective is to understand the hierarchical viewpoint of security that each represents.
- ▶ Apply your view of each standard to products that you use on your desktop everyday. Could it meet some level of the standards? What configuration changes would you need to do, in your estimation, to meet the standards?
- ▶ Determine whether the product you use on your desktop has been evaluated. At what level has it be evaluated—at a granular level or in general? When was it last evaluated? If it hasn't been evaluated, can you determine why it might fail or succeed if it were to be evaluated?

“The security architecture and models domain contains the concepts, principles, structures, and standards used to design, implement, monitor, and secure operating systems, equipment, networks, applications, and those controls used to enforce various levels of confidentiality, integrity, and availability.

The candidate should understand security models in terms of confidentiality, integrity, information flow, commercial vs. government requirements; system models in terms of the Common Criteria, international (ITSEC), United States Department of Defense (TCSEC) and Internet (IETF IPSEC); technical platforms in terms of hardware, firmware, and software; and systems security techniques in terms of preventative, detective, and corrective controls.”

—Common Body of Knowledge study guide

This chapter covers Domain 6, Security Architecture and Models, 1 of 10 domains of the Common Body of Knowledge (CBK) covered in the Certified Information Systems Security Professional Examination. I have divided this domain into several objectives for study.

INTRODUCTION

How do you say security? Today it’s popular to speak of it, but I don’t think most people have learned to pronounce the word yet.

Perhaps it’s the manager, CIO, someone with the purse strings who will approve anything with the word *security* in it. Firewall. Yeah, give me one of those. Intrusion detection, PKI, smart cards, and tokens—I’ve got lots of security right here, folks.

Or maybe she’s a network administrator. Can’t wait to play with these new toys? Or she’s found security to interfere with performance. Until management changes the directives, security is just another thing to keep running and keep out of the way of getting data from here to there—fast.

Could be he’s a programmer, or project manager. Security? Why he’ll build that right into the product. Crypto, access controls, public keys, no worries. They say it’s hard to get it right? Bring it on.

Then there's Joe. Hi, Joe. Joe just wants to get his job done. He doesn't want to configure a personal firewall, select a secure operating system, or learn anything new. But Joe doesn't want his identity or money stolen.

And maybe, just maybe, the people with the power to institute sound info-security practices realize the previous reactions for what they are. Perhaps you're one of them. If so, how does anyone build, buy, and use more secure products? How do you make your infrastructure more secure?

Here's how. You find out about the joint efforts of those who came before you and what they have said about it. You look for the standards, the validated practices, and certified products that are out there. No one person has the answer. There is always much to learn, and research is continual; but you don't have to do it alone or rely on commissioning your own study from the ground up. There is a tremendous amount of information available. I'm not talking about academic research, I'm talking about real-world implementable designs that have been and are being used by governments, by financial institutions, by utilities, commercial industries, and organizations around the world. I'm trying to point you to products that have been evaluated against these programs, hoping you will use this information to build or improve your own security operations. That's what this chapter is about—architecting security. Taking the models, the schemas for secure products, the assurance formulas that exist, and applying them to a real-world environment. If this is already your modus operandi, no offense meant, but if you like the way I talk about it, please pass it on.

All these things are important, but you just can no longer expect to bandage your systems with security products, which mask your fragility by creating born-again security awareness from software developers. Here's my point. You've got to architect your information systems like they were meant to stand up to more than script kiddies and virus-writer wannabes. It is not a problem you can throw people, or product or money at. Instead, it's a constant, all-encompassing movement. This chapter introduces you to some of the work that has gone on before. Pick up the flame and run with it.

NOTE

Computer Trustworthiness = Trustworthy Computing? Study computer security long enough and you'll stumble across the concept of computer trustworthiness. That is, a computer is trustworthy if it has a trusted computing base, enforces a security policy, and has domain separation, resource isolation, hardware isolation, software isolation, and software mediation. This "trustworthy" characteristic of a computer system sounds like a component that's needed in "trustworthy computing," an initiative that Microsoft has pledged to work for; visit the progress for this project at <http://www.microsoft.com/mscorp/execmail/2002/07-18twc-print.asp>. How do their products, and those that you use, stand up?

REQUIREMENTS FOR SECURITY ARCHITECTURE AND MODELS

Explain the difference, if any, between public versus government requirements for security architecture and models.

Historically, government computer security issues have centered on confidentiality—making sure unauthorized individuals cannot access information. On the public, or commercial side, concerns have been of the correctness or integrity/consistency of data. The security models—Bell LaPadula (a government access control model that addresses confidentiality) and Clark-Wilson (written for commercial concerns and addresses integrity), both described in the following sections—seek to address these concerns; and the earliest security architecture, the Orange Book (government sponsored and mainly concerned with confidentiality), does as well.

A second difference has been the tendency to consider governmental information as requiring much more security against theft or manipulation. After all, the exposure of confidential commercial data might cause a business to fail. The exposure of government information might topple a state.

Third, only the very largest and wealthiest businesses saw the need for, or could afford to apply, information security practices and products routinely used in governmental affairs.

Both government and public concerns have data of varying sensitivity, and both have used a variety of techniques to vary the level of security applied to different data classifications. Governments may use classifications such as *unclassified*, *classified*, *secret*, *top secret*, and *eyes only*, whereas businesses generally use the terms *public*, *private*, and *confidential*. Many government records are public—that is, they are available to anyone with the wherewithal. In the past, that has meant the ability to physically locate and spend time searching through microfiche and ledgers, or to file numerous documents and pay copying fees in order to obtain them. Now it might mean downloading them from the Internet. Businesses also have public information, product data, advertising, and so on, which are visible to the public.

Governments and businesses have more sensitive information that is kept confidential: troop movements, top-secret research, employee salaries, financial data, trade secrets, and so on.

Neither governments nor commercial entities envisioned the explosion of communication and interconnectivity fostered by the growth of the Internet and the ubiquity of computing. Many things have changed. For example, consider the following:

- ◆ The average teenager in America can purchase, and may already own, computing power and connectivity unavailable to the most sensitive government offices, or the richest commercial enterprises just a decade ago. Small, poor nations and non-affiliated terrorists cells can own computing resources that are adequate enough to attack any business, government, or infrastructure anywhere in the world. Prewritten scripts and write-your-own-virus engines exist and can be freely downloaded off the Internet. Wireless connectivity allows access across the traditional barriers of cable and connection. From anywhere in the world, these resources can be used to attack, disrupt, and compromise almost any computing system almost any place.
- ◆ Even if a business chooses to disregard these security threats as minimal, they must consider the easy familiarity most people have with computers today. Less than ten years ago it was quite common to find employees who were petrified of computers. Now many, if not most employees have had many opportunities for computer use and for classroom training.
- ◆ Improperly configured systems expose data of all kinds to accidental manipulation and misuse. Years ago, few were sophisticated enough in computer technology to take advantage of this fact, and many would say, fewer systems were so simplistic as to make it so easy to do.

For these reasons and more, there is less and less difference between the needs of government and public enterprise for security models and architecture. The process is the same. The threats must be understood, the risk analyzed, the products researched, and the plan developed.

SECURITY MODELS

Discuss examples of security models including the following:

- **Bell-LaPadula**
- **Biba**
- **Clark-Wilson**
- **Access control lists**

NOTE**The ICS2 Approach to These Models**

You'll find that these models were also discussed in other areas of the book. The ICS2 is redundant regarding the information covered in each domain. In some cases, this is due to the context in which these models are discussed. In other cases, it's redundancy for the sake of redundancy. Obviously, it's important that you have an understanding of each model, thus, we once again approach these models with the ICS2 domains in mind.

A security model is a prescriptive paradigm. At first, it's someone's best guess at formulating a plan to make something more secure. It gets tested, refined, used, and maybe abandoned as the "things" you're trying to secure and the resources you have to do so change. Nevertheless it is important to know about them. They may be in place where you work, or they may lead you to a better understanding of your job. Their study will also teach you the vocabulary of modeling secure systems. The following security models are a few of the better-known ones:

- ◆ Bell-LaPadula
- ◆ Biba
- ◆ Clark-Wilson
- ◆ Access control lists

Each of these is discussed in the following sections.

Bell-LaPadula

Bell-LaPadula is an information flow security model. This model was developed in the 1970s in response to the U.S. government's concern about security on the mainframe systems on which it used. The main issue was confidentiality, how to keep unauthorized personnel from accessing data. Access to stored data could be controlled through access controls that identified who could access what. But, what happens when data is moved? The Bell-LaPadula model has as its premise that "information shall not flow to an object of lesser or non-comparable classification." To understand what is meant by that I'll detour into some basic security modeling explanations.

Two key terms you need know are *object* and *subject*. By object, I mean passive items such as hardware, software, and processes that store information. The *subject*, however, is used to describe active processes, such as persons or devices that move information between objects. Each subject, even if it acts on behalf of another subject, is assigned a formal security level or clearance. Each object is also given a security level or classification. Object and subject security levels are identified by assigned labels.

An easy example of this object-subject relationship is to think of the nature of government, business, or even personal information. Let's use a publicly traded business example. For this business, some information is public knowledge. Names, addresses, contact numbers, and other quarterly information about the stocks are public information. Other information, such as day-to-day financial transactions, is for only those processing the transactions, and certain management personnel. Still other information (the financial health and well being of the company before the public announcements or going beyond what is appropriate and legally obligated in those statements) is severely restricted. For this business, as for your personal life, data has different classifications and the ability to access information is controlled. You don't have to formally label it classified, unclassified, secret, top-secret, or eyes only, in order for it to be so. Conversely, we give individuals within our sphere of influence (business, personal life) different levels of clearance to see our information. (Your lawyer, for example has much more privileged information about you than I do.)

So, loosely translated, Bell-LaPadula is saying that one of the ways data can be kept secure is if the data is never moved from a container classified at level X to another container that has a classification lower than X, or that cannot be judged to be of equal or higher classification. Practically speaking, it's as if you agreed to keep your cash safe by never moving it from the bank vault to your wallet. You can move it to another bank vault, but not to the wallet, pocket, hand, or refrigerator. Note that I'm not talking about being able to practically use that data (money in the last example); I'm merely talking about how to keep it safe.

Why does this security model work? It works because it presumes (and explains) that access to each classified container, or object, is also strictly controlled. That is, every subject must have clearance; they must be authorized to access the container. It also eliminates possible covert channels (ways of communicating information without seeming to do so).

One classic covert channel might exist in some systems because you strongly protect access to objects and take great pains to selectively grant the rights of access to these objects, but fail to prevent the movement of data from one object to another. Picture, for example, the results if I have authorization to read and write to the file directories A and B. Folder A has personnel records in it and is on a computer drive where access permissions can be set. Folder B has a document detailing the weekly lunch menu in the company cafeteria. Folder B is on a computer or drive where access permissions cannot be set. Access is controlled by permissions on its entry. You do not have access to folder A, but as an employee of the company you have read access to the lunch menu folder, folder B. Because I have read and write access to both folders I can copy the personnel records, which of course, include salary information, from folder A to folder B. Now others can read them too!

If our system followed the Bell-LaPadula model, I would not be able to transfer the personnel files to a publicly available folder. Extending this concept, any subject that has authorization to access A-level data does not have write access authorization to B-level data. Other subjects may have write access to level-B data, and they may have less clearance than those with access to A. The danger of transfer of information to an object of lower classification is prevented.

This has been a much-simplified description of this model. The model itself has much more to it. One of its premises is that it follows the computer science Basic Security Theorem. This theorem states that a system can be put into a secure state that is security preserving. That is, a sequence of rules applied to the system in a secure state will result in the system entering a new secure state. The theorem, and the Bell-LaPadula model can be proven using set theory and other mathematics. Some other basic concepts of Bell-LaPadula are

- ◆ **Fundamental modes of access**—Access, such as read, write, read only, and so on, is defined to permit access between subjects and objects.
- ◆ **Dominance relations**—A relationship between the formal security levels of subjects and objects describes the access permitted between them.
- ◆ **Simple security condition**—A single statement such as granting read access to a specific object. For example, “Grant Bob read access to file B.”

- ◆ **Discretionary security property**—A specific subject is authorized for a particular mode of access that is required for state transition.
- ◆ **Star * property**—information cannot be written to another lower level.
- ◆ **Trusted subject**—Access under this option is not constrained by the star * property.
- ◆ **Untrusted subjects**—Access under this option is constrained by the star property.

Biba

Where Bell-LaPadula address secure information flow and confidentiality, the Biba model was the first to address integrity in computer systems. In this model, no subject may depend on a less trusted subject, and the primary objective is to prevent users from making modifications that they are unauthorized to do.

Biba is based on a hierarchical lattice of integrity levels and is an information-flow security model. In this model, two rules prevail—no *write up* and no *read down*.

First, no subject can write up to a higher integrity level. Let's think about the request I might make at the bank for some money. I make out a check for \$100.00 to "cash" and hand it to the teller. I'm telling her that I have \$100.00 in my checking account and I would like her to give it to me. The teller, however, does not take my word about the data in my account (the lower integrity level) for truth. Instead, she checks the bank's computer records (the higher integrity level). If the funds do indeed exist, she gives me the cash and the work is started to reduce by \$100, the balance of my account.

Second, no subject can read down. In our example, the bank computer does not need to read any balance information from the request that I make. The teller may enter the information that I am withdrawing \$100, and even the balance left in my account, but the transaction that records the information will not use this information, and the processes that manage the account balances have no authority to read the file that contains the information.

Clark-Wilson Model

The Clark-Wilson Model also emphasizes data integrity, and does so for commercial activities. It uses software engineering concepts such as abstract data types, separation of privilege, allocation of least privilege, and non-discretionary access control. Clark-Wilson has three integrity goals:

- ◆ Prevent unauthorized users from making modifications
- ◆ Prevent authorized users from making improper modifications
- ◆ Maintain internal and external consistency

Much of the implementation of this model consists of using well-formed transactions that preserve consistency—for example, the user can only modify data in ways that ensure internal consistency. If I visit my branch of the bank in which I have my checking account and transfer money from savings to checking, I would be pretty upset if the money was removed from my savings account, but somehow never showed up in my checking account.

What looks like a simple action to me (take from B and put into C) to the bank's database is a two-step transaction. First the account balance of the savings account is reduced by the amount I specified, and then the balance of the checking account is increased. Because this does not, and cannot, happen at the same time, it might be possible that, say, if the computer crashed in between these two operations, my checking account balance might never be increased.

Fortunately for me, this particular problem of database consistency has been solved for a long time. Modern database management systems track the completeness of transactions and ensure that should something happen, say the computer crash as mentioned previously, the transaction is either rolled back (the money is returned to my saving account) or rolled back and then reapplied as it was meant to be.

Clark-Wilson prescribes this philosophy to all possible data modifications to ensure integrity and consistency.

Access Control Lists

Unlike the formal security models described previously, the access control list security model is familiar to a wide population of IT people. Some of these people are system and network administrators in Unix and Windows environments, or desktop users with Windows NT Workstation, or Windows 2000/XP Professional. In addition, IT managers, IT auditors, and others have learned the model from first-hand experience of its implementation.

In this model, objects (the resources) are assigned lists of approved subjects (users and groups). Each entry in the list consists of user identification of some form, and the approved access level. Access levels are appropriate for the resource—hence for files, levels may be read, write, read/write, and so on, whereas for printers, levels may be manage or print. Subjects, the users and groups, are assigned some kind of identification.

A security kernel, or reference monitor, serves as the arbitrator of access requests. The subject's request must match his identification and authorization as listed in the object's access control list, or he is refused.

This is an effective, flexible system, but it has the potential for complexity and confusion. What, for example, is the result of saving a new file in a directory with a certain access control list? Are the lists inherited? Are they modified by special characteristics? And what happens if inheritance is the rule and changes are made to the top of multiple layers of directories? Do multiple permissions set on varied resources affect the outcome?

Furthermore, unlike many labeling systems, access control lists can be rapidly changed, resulting in different and unexpected behavior.

REVIEW BREAK

A Review of the Security Models

Four security models, all of which apply to access control, have been discussed here. Table 6.1 summarizes them.

TABLE 6.1
SECURITY MODELS FOR ACCESS CONTROL

	<i>Government Model</i>	<i>Primary Directive</i>
Biba	Yes	Integrity
Bell-LaPadula	Yes	Confidentiality
Clark-Wilson	Yes	Integrity
Access control lists	No	Attempts at both confidentiality and integrity but limited to proper application

SECURITY SYSTEM ARCHITECTURE

Explain the basics of security architecture.

A security architecture is the sum of the components used and the way they are put together to build security functionality into a computer operating system, device, or system. Many make the mistake in qualifying a system as either secure or non-secure, when in reality a wide range of security features and functionality may be designed into a system. In addition, most modern systems can operate in different modes, either through selection or misconfiguration. Therefore, selection of a secure system must go beyond the binary to a matching of need against delivery and an understanding of configuration versus accreditation. This section introduces terminology and concepts that are useful in understanding any discussion of system security.

Reference Monitor

One of the primary concerns in the evaluation of the security of systems is how the system controls access. Does it use labels or permissions? Are controls mandatory or discretionary? How granular is it? Is there any way around it? A key component in any secure system implementation is the one that controls this function, the reference monitor. The reference monitor is an imaginary device that controls all access to all objects (passive items such as hardware, software, and processes that hold or store information) by subjects (active processes, persons, or devices that move information between objects).

Think of the reference monitor as if it were some internal security control center for a building with many doors. Access to resources behind each door can be requested by using the phone at the side of the door. The phone connects you to the security control center. The security control center checks your credentials and your request. If they match a list that identifies you as someone who can access that door, the door opens. If you aren't identified as someone who can access it, the door does not open. You may approach another door, repeat the process, and be allowed in—assuming your credentials are verified. The control center, or reference monitor, has done its job. Each attempt at access is carefully screened, and access is granted or denied.

Those familiar with the security subsystem of Windows NT and above will recognize the component called the Security Reference Monitor (SRM). The SRM, which examines the credentials of the requestor for access to resources (or objects, such as files, registry keys, and printers), either permits or denies the request. Figure 6.1 illustrates the concept. Windows NT uses a Security ID or SID to identify subjects (subjects are user accounts and security groups). Objects are assigned Access Control Lists (ACL) which consist of entries identifying by SID the type of access (read, write, and so on) a subject may have. When a subject logs on, a list of his credentials (SIDs for ID and group membership) is compiled and placed in an access token.

1. JohnS attempts to access the "accountants" folder.
2. The Security Reference Monitor (SRM) checks security identifiers (SIDs) in John's access token against security identifiers in the Access Control List (ACL) for the "accountants" folder.
3. If the SRM finds a match between the token and the user's request, and the ACL, access is allowed. If not, access is denied.




FIGURE 6.1

Requests are channeled through the Security Reference Monitor, which matches the SIDs of the subject's access token against the list of SIDs and permissions associated with the object. If no match is made, access is denied.

Open Versus Closed Systems

In early computing history, just being able to harness the power for mathematical use was enough. As systems evolved and the nature of the tasks they were used for became more diverse, the need to ensure the confidentiality of information stored on them also grew. Early users, many of them governments and military operations, required a system that could be secured against unauthorized use. Systems were designed and built to be either secure or not secure. This is where the concept of open versus closed systems was developed, and where some of our problems exist today. By definition, an open system provides a user with total systems access—in effect, he is the administrator of his machine. A secure system, on the other hand, is totally secure.

The problem with this idea of secure or non-secure tricks people into believing that a computer that they have used and that has security features is secure, while another that they have not used, is not secure. It allows the unsophisticated to confuse an off-the-shelf production model with a hardened (configured for security) system. Another myth is that there is a clear way to categorize specific systems as either secure or not secure. In reality, it is extremely difficult to produce a system for today's needs that can be in itself 100% secure. What you can do is design and build systems that can be secured, and even so you must continually maintain them in order to keep them so. At the other end of the spectrum, it is possible to produce a system with no security controls whatsoever. Table 6.2 compares the features of open and closed systems.

TABLE 6.2

AN OPEN SYSTEM VERSUS A CLOSED SYSTEM

	<i>Open</i>	<i>Closed</i>
User interface	Standard	Nonstandard
User access to system	Total	Limited to a single application or language

By definition then, a large number of computer systems are open systems, whereas few are closed. However, many of the open systems now have security features that can be configured to make them more secure.

Systems such as many modern Unix systems and more recent versions of Windows systems (Windows NT, 2000, XP, and .NET) default to a single administrative account, and provide the ability to create users who are limited in their privileges on the system. In addition, these systems provide discretionary resource access control. These systems are not, however, closed systems, though granular control of user access can contain a user to a single application.

In sum, while there is still a need to distinguish between open and secure systems, you should be careful not to assume that all systems are either one or the other and that even the most secure system must be configured to be so, and must be maintained to stay secure.

Security Principles

A good security system architecture is designed to maximize the use of recognized security principles. Among these principles are

- ◆ **Trusted Computing Base (TCB)**—The sum of the security functions of the system.
- ◆ **Execution domain**—The OS system area is protected from tampering and accidental modification. In many systems this is implemented by creating a secure area, or kernel, within which the operating system functions. Another layer, the user area, is set aside for application programs.
- ◆ **Layering**—Processes do not do everything. Processes are layered, with each layer having a specific job. An example of this functionality is the requirement for user applications running in the user area of the system, to call kernel-level functions when necessary access to system operations is required.
- ◆ **Abstraction**—Acceptable operations are characterized, not spelled out in detail.
- ◆ **Process isolation**—Many processes can be running without interfering with each other. In many systems this means each process is assigned its own memory space.
- ◆ **Least privilege**—A process has only the rights and access it needs to run; only processes which need complete privileges run in the kernel and other processes call on these privileged processes only as needed.

- ◆ **Resource access control**—Access to resources is limited.
- ◆ **Security perimeter**—The boundary of the TCB. A security kernel and other security-realized functions operate within this perimeter. A security kernel is the implementation of the reference monitor concept.
- ◆ **Security policy enforcement**—The policy set for the system must be operational in order for the system to be operational, the security policy is always followed.
- ◆ **Domain separation**—The objects that a subject can access become its domain. For example, users generally have access to run programs and open and write to certain files. The user doesn't need to access the security kernel, for example, so the domain of the TCB is separated from that of the user.
- ◆ **Resource isolation**—Subjects and objects are kept separate for control purposes.

Security Modes

A security subsystem may be designed to operate in a particular mode. The mode is based on the need to authorize access to different levels of data sensitivity. This is one way to view both the nature of the data available on the computer, and the restrictions on access. The modes are

- ◆ **Dedicated**—No restrictions. All users can access all data. All users have clearance for all data on the system and have signed nondisclosure agreements for all information stored and processed. The users have a valid need to know for all information.
- ◆ **System high**—All users have access approval and clearance for all information on the system. Users have clearance for all information, they have a need to know for some of the information, and signed nondisclosure agreements that require them not to share the information.
- ◆ **Compartmented**—Users have valid clearance for most restricted information processed on the system, formal access and non disclosure for that information, need to know for that information. Data is partitioned. Each area of data has different requirements for access. Users of the system must meet the requirement for the area they wish to access.

- ◆ **Multilevel secure (MLS)**—Users have different levels of clearance to different levels of information (think Bell-LaPadula). Some do not have valid personnel clearance for all information; all have valid need to know for that info to which they have access
- ◆ **Controlled mode**—Multilevel in which more limited amount of trust is placed in the hardware/software base of system. This results in more restriction on classification level and clearance levels.
- ◆ **Limited access mode**—Minimum user clearance is not cleared, and maximum data sensitivity is not classified by sensitivity.

Labels Versus Access Control Lists

The earlier discussion of security models and many discussions of security systems mention labels and labeling as a system for use in access control. However, many modern, commercial computer systems use access control lists instead. Which is better? There is no easy answer here.

Using labels presents the opportunity for more rigid control. For example, a user may be only permitted to initialize sessions with a specific label. Since labels for resources in many systems, once set, cannot be changed, it is possible to predict with a fair amount of certainty what the user will be able to access. Access control lists, on the other hand, can be modified; systems using them often allow user sessions to be restricted only by the ability to match their credentials to the access control lists. User credentials can also be changed. Well-managed systems that apply other security controls (the rights to modify access controls and user credentials is restricted, the decision to modify them is controlled by policy and the policy is enforced) can maintain access control as set by policy. Poorly managed systems that allow arbitrary changes to occur make this an impossible chore.

Using labels would seem to offer the opportunity to make systems more secure. On the other hand, they are usually very expensive to administer, and their rigidity makes them difficult to use in a complex world of shifting requirements.

Covert Channel

It is important to understand the concept of a covert channel because it is often an unexpected vulnerability in an otherwise secure and securely maintained system. Being able to recognize such a flaw may lead to its prevention.

A covert channel allows an object with legitimate access to information to transfer the information in a manner that violates system security policy. Two types of covert channels exist—covert storage channels and covert timing channels.

The covert storage channel allows the direct or indirect writing by one process to a storage area that allows direct or indirect reading by another process which has less clearance than the first. In essence, it's as if an individual with security clearance leaves top-secret information lying around on a table at the food court in a mall. This is similar to when a disk space is shared by two objects that have a different security classification. In a simple labeling system, subjects with clearance for either classification have access to the disk. In an access control list protected system, a folder (directory) has permissions set that allow both subjects access. When the more sensitive information is saved on the disk, both sets of controls are applied, and either subject can access the files.

A covert timing channel exists when a signal of information is modified due to some other system function. The modified signal may allow unauthorized individuals to determine the system function through observation of the other. For example, a recent study concluded that the disk access lights on a system, when carefully studied, reveal information about the data being processed on the system.

While covert channels are often the result of system design or configuration, an exploitable channel is a covert channel that is created with the intention of violating security policy. It is useable or detectable by a subject external to the trusted computing base.

INFORMATION SYSTEM SECURITY STANDARDS

Describe information system standards, including the following:

- **TCSEC**
- **ITSEC**
- **Common Criteria**

When information security requirements are high, an evaluation of computing systems and devices should be done before the systems are put into production. If formalized, as is required in many government operations, this process consists of two steps. First, the system is given a technical evaluation and is certified to have the security features that are specified for the job for which it will be used. Second, management must decide to accept the risk of using this system and approve its operation and environment. The management evaluation may result in approval (accreditation) or rejection. In addition, if the systems are to be configured to meet the evaluated circumstances, the objective may be to have the site certified. This type of accreditation requires outside authority and is beyond the scope of mere administrator configuration and local management approval.

The diverse nature of computing needs, as well as the capability of computing systems to fulfill them, can create a backlog of requests if each new product must be technically evaluated. Early efforts in the United States to resolve this issue resulted in the Trusted Computer System Evaluation Criteria (TCSEC)—a U.S. Department of Defense standard for computer system security. Better known as the Orange Book due to the color of its cover, this standard consists of a rating system against which systems could be formally evaluated. The receipt of a rating relieved an individual government department from doing the lengthy technical evaluation on its own and prevented duplication of efforts.

NOTE

Standards to Know Historically, several security evaluation systems are of note:

- Orange Book—Trusted Computer System Evaluation Criteria (us) (TCSEC)—1985
- UK Confidence Levels 1989
- ITSEC (1991) Information Technology Security Evaluation Criteria (from the German and French Criteria, and the Netherlands, and United Kingdom)
- Canadian Criteria 1993 Canadian Trusted Computer Product Evaluation Criteria (CTCPEC), a combination of ITSEC and TCSEC
- Federal Criteria 1993 (draft Federal Criteria for Information Technology Security)—later merges into Common Criteria

NOTE

What Meaning Can You Ascribe to the Security Rating? In the United States, the Orange Book certifications have long been perceived as indicators of securable systems to those who understood the standard, and of secure systems to those who didn't. That is, the very existence of a rating was often interpreted as meaning the system was secure. Instead, the certification only means that that particular version of the system is securable to that level, when configured precisely as the system tested, and running on the hardware and operating environment specified in the evaluation. The easiest case in point is the C2 rating received for Windows NT 3.51. Windows NT 3.51 does not install to C2 certification specifications. Care must be taken to apply the evaluation criteria, and, if properly done, the system will be unable to function in the manner in which most purchasers want it to function. This is not a complaint against Windows NT 3.51, but against those who do not understand what having a certification means.

Other governments developed additional standards that improved on this early one. Most notable are the European model, the ITSEC which is accepted by several European nations, and the Canadian standard, CTCPEC. More recently, a number of governments merged these existing standards into the Common Criteria, an internationally recognized information system security standard. Commercial enterprises have also used these systems and have been instrumental in their continued evolution. In the United States, Common Criteria is the standard being used; however, many existing systems have ratings from the earlier Orange Book. In Europe, many countries have accepted the Common Criteria standard, but also have existing ratings in use.

It is important, therefore, to understand something of the provisions of the earlier evaluation systems to help us understand the current system, and because they are still in use today. These guides describe, in sometimes excruciating detail, the specifications against which the products should be judged.

For all of them, third-party evaluation and certification is the ultimate goal. If a product received a particular certification level, that product would be accepted by the entities that accepted the standards.

You should be aware of the following standards because products certified by all of them are available in the marketplace:

- ◆ TSCEC
- ◆ ITSEC
- ◆ Common Criteria

TCSEC—The Orange Book and the Rainbow Series

The certification emphasis of the Orange Book is confidentiality. The concept of a secure, or trusted, system is divided into a series of classifications that range from minimal protection to verified protection. As the use of the system continued, a series of additional guides were written to support its use and to describe the implementation of security principles that were not addressed in the original guide. This series of books is referred to as the Rainbow Series, and each book is also identified by the color of its cover. You can download copies of each of these guides, and the Orange Book, at <http://www.radium.ncsc.mil/tpep/library/rainbow/>.

Orange Book Classifications

The Orange Book outlines the evaluation criteria and gives an objective measure for acquisition. It divides operating systems into four primary divisions around three different concepts. The concepts are

- ◆ Ability to separate users and data
- ◆ Granularity of access control
- ◆ Trust or overall assurance of the system

The primary divisions are

- ◆ **D**—Minimal protection
- ◆ **C**—Discretionary protection
- ◆ **B**—Mandatory protection
- ◆ **A**—Verified protection

The primary divisions are further divided into classes, as described in Table 6.3. Within each class, evaluation is based on six fundamental security requirements and the system documentation. They are

- ◆ **Security policy**—Must be explicit and defined by the system.
- ◆ **Security policy**—Must include some form of marking; access control labels must be associated with objects.
- ◆ **Accountability**—This is ensured by requiring the identification of all subjects.
- ◆ **Accountability**—Determined by being able to audit information and attribute actions to individuals.
- ◆ **Assurance**—This is possible by using evaluated hardware and software that enforces security policy.
- ◆ **Continuous protection**—This is ensured because trusted mechanisms protect the system and are themselves protected against tampering and unauthorized changes.

NOTE

It's Not Perfect Security architecture models address operating system security. They address system access controls, data access controls, system security, and administration and system design. They do not address the issues of physical security nor do they deal with the human factor.

TABLE 6.3**ORANGE BOOK CLASSES**

<i>Class</i>	<i>Title</i>	<i>Description</i>
D	Minimal protection	Have been evaluated but don't meet standards for other classes.
C	Discretionary protection	Need to know protection, accountability of subjects, accountability of actions, and audit.
C1	Discretionary security protection	Separation of users and data, enforces access limitations, users use data at the same level of security.
C2	Controlled access protection	More granular, user is more individually accountable, logical procedures, auditing, resource isolation; security policy enforcement, accountability, assurance. Controls who can log in, access to resources based on wishes or users, log of user actions.
B	Mandatory protection	Integration of sensitivity labels, labels used to enforce mandatory access rules, specification of TCB, reference monitor concept implemented.
B1	Labeled security protection	Accurate labeling of exported information.
B2	Structured protection	Formal security model, discretionary and mandatory access control extended to all subjects and objects. Covert channels addressed. TCB has protection critical and non-protection critical elements, trusted facility management (systems admins and operator functions, configuration management control). System is relatively resistant to penetration.
B3	Security domains	Reference monitor must mediate all access of subjects by objects, and is tamperproof. Unauthorized code is excluded, security policy enforcement, complexity minimized, security administrator supported, audit expanded, and system recovery are required. System is highly resistant to tampering.
A	Verified Protection	
A1	Verified design	Functionally equivalent to B3, but verification techniques are used against the formal security policy. Can give high degree of assurance. TCB is correctly implemented.

Criticisms of Orange Book

Several criticisms of this evaluation system exist. Following are the major ones:

- ◆ The Orange Book criteria primarily address confidentiality, or the concept that if you control how users get to information, you don't have to worry about correctness of data. Unfortunately, that is not always the case. Banks and many others want assurance that data is correct.
- ◆ In addition, the Orange Book emphasizes controlling users, but doesn't say anything about what users might do with the information they get.

- ◆ It does not fully address procedural, physical, and personnel safeguards, nor how the safeguards might impact system security.
- ◆ It does not address networked computers. (The later published Red Book of the rainbow series does this.)

Although these criticisms are correct and are the reason that the newer, international standard, Common Criteria, is now accepted, you should always remember the climate and status of computing at the time when this system was developed. It was developed at a time when computing consisted primarily of mainframe systems used by government installations and extremely large commercial enterprises. It was developed by the United States Department of Defense (DoD) and so primarily addressed needs defined by the DoD. Additional guides in the Rainbow Series address many of these criticisms.

Rainbow Series

There are some 30 security guides that supplement or explain the Orange Book. Each book is referred to by the color of its cover. (There is no significance to the color.) One of the more important of these guides is the Red Book. This book interprets the TCSEC in terms of networking. Some other examples of interpretations in the Rainbow Series are described in Table 6.4.

TABLE 6.4

OTHER INTERPRETATIONS IN THE RAINBOW SERIES

<i>Number</i>	<i>Title</i>	<i>Common Title</i>	<i>Publication Date</i>
CSC-STD-002-85	<i>DoD Password Management Guideline</i>	Green Book	4/12/85
CSC-STD-003-85	<i>Computer Security Requirements 0 Guidance for Applying the DoD TCSEC in Specific Environments</i>	Light Yellow Book	6/25/85
CSC-STD-004-85	<i>Technical Rational Behind 003-85 (above)</i>	Yellow Book	6/25/85
NCSC-TG-001 Ver 2	<i>A Guide to Understanding Audit in Trusted Systems</i>	Tan Book	6/1/88
NCSC-TG-002	<i>Trusted Product Evaluations—A Guide for Vendors</i>	Bright Blue Book	6/22/90
NCSC-TG-003	<i>A Guide to Understanding Discretionary Access Control in Trusted Systems</i>	Neon Orange Book	6/30/87
NCSC-TG-004	<i>Glossary of Computer Security Terms</i>	Teal Green Book	10/21/88

continues

TABLE 6.4*continued***OTHER INTERPRETATIONS IN THE RAINBOW SERIES**

<i>Number</i>	<i>Title</i>	<i>Common Title</i>	<i>Publication Date</i>
NCSC-TG-005	<i>Trusted Network Interpretation of the TSCEC</i>	Red Book	7/31/87
NCSC-TG-006	<i>A Guide to Understanding Configuration Management in Trusted Systems</i>	Amber Book	3/28/88
NCSC-TG-007	<i>A Guide to Understanding Design Documentation in Trusted Systems</i>	Burgundy Book	10/06/88
NCSC-TG-008	<i>A Guide to Understanding Trusted Distribution in Trusted Systems</i>	Dark Lavender Book	12/15/88
NCSC-TG-009	<i>Computer Security Subsystem Interpretation of the TCSEC</i>	Venice Blue Book	9/16/88

Information Technology Security Evaluation Criteria

This European standard was developed in 1991 by Germany, France, the Netherlands, and the United Kingdom. In 1998, Finland, France, Germany, Greece, Italy, Netherlands, Norway, Spain, Sweden, Switzerland, and the United Kingdom agreed to recognize Information Technology Security Evaluation Criteria (ITSEC) certificates from Qualifying Certification Bodies—for example, *Serveur thématique sur la sécurité des systèmes d'information* (SCSSI; <http://www.scssi.gouv.fr/fr/index.html>) of France, *Bundesamt für Sicherheit in der Informationstechnik* (BSI; <http://www.cert.dfn.de/eng/csir/europe/bsicert.html>) of Germany, and *Communications-Electronics Security Group* (CESG; <http://www.cesg.gov.uk/>) of the U.K.

In 1999, these countries, with the exception of Germany, agreed to accept Common Criteria up to the EAL7 level. (Germany accepted Common Criteria evaluations in 1998 to EAL4 level.) These countries, like the United States, are adopting the Common Criteria and any discussion of ITSEC is often tempered by comparison to Common Criteria. In the United Kingdom, the official ITSEC Web site <http://www.itsec.gov.uk/> has been subsumed by the newer Assurance site <http://www.cesg.gov.uk/assurance/iacs/itsec/index.htm>, which is the official page now. This site also speaks to the adoption of Common Criteria and provides comparative information.

Differences Between the Orange Book and ITSEC

Several differences exist between ITSEC and the Orange Book:

- ◆ Unlike the Orange Book, which concentrates on confidentiality, ITSEC addresses the triple threat of loss of confidentiality, loss of integrity, and loss of availability. Those familiar with information security dictums will recognize the famous CIA (confidentiality, integrity, and availability) triad.
- ◆ In the specifications, the Target of Evaluation (TOE) is the product or system to be evaluated. The TOE's functionality (can it provide this security function) and Assurance (how do you know it is providing this functionality) are evaluated separately.
- ◆ ITSEC does not require the security components of a system to be isolated into a TCB.
- ◆ ITSEC provides for the maintenance of TOE evaluation. Some systems can maintain certification after patches, without formal reevaluation.

The separation of functionality and assurance is accomplished by recognizing three objectives of evaluation:

- ◆ **Security functions**—What is done.
- ◆ **Security mechanisms**—How it is done.
- ◆ **Certification**—The TOE meets the security target to the claimed assurance level.

The United Kingdom Information Technology Security Evaluation and Certification Scheme

Like the Orange Book, the ITSEC levels of certification are scaled; each level includes increasing security functionality. Certification is carried out by Commercial Evaluation Facilities or CLEFs, which are appointed by the Certification Body of the Scheme. Table 6.5 lists and describes the levels.

NOTE

Certification A certification is a formal statement confirming the results of an evaluation and confirming that evaluation criteria were correctly applied. The evaluation is conditional and is only true when the TOE is configured and used in the manner in which it was evaluated. Certification does not endorse the TOE, nor guarantee its freedom from exploitable vulnerabilities.

NOTE

CIA Computer security is often defined as the combination of these three principles: confidentiality, or the prevention of unauthorized disclosure of information; integrity, the prevention of unauthorized modification of information; and availability, the prevention of unauthorized withholding of information or resources.

TABLE 6.5**ITSEC LEVELS**

<i>Level</i>	<i>Description</i>
EO	Inadequate.
E1	Definition of security target and informal architecture design exists, User/Admin documentation on TOE security. TOE is uniquely identified and documentation exists which includes delivery, configuration, start-up, and operations. The evaluator tests the security functions. Secure distribution methods are utilized.
E2	Informal detailed design and test documentation are produced. Separation of TOE into security enforcing and other components. Audit trail of start up and output required. Assessment includes configuration control, developer's security and penetration testing for errors.
E3	Source code or hardware drawings must accompany the product, and a correspondence between design and source code must be shown. Standard, recognized implementation languages are used. Retesting is required after correction for errors.
E4	Formal security model. Semi-formal specification for security enforcing functions, architecture, detailed design. Sufficient testing. TOE and tools under configuration control. Changes are audited, compiler options documented. TOE retains security after a restart from failure.
E5	Relationships between security enforcing components are defined in architectural design. Integration processes and runtime libraries are provided. Configuration control is possible independently of developer. Configured, security enforcing or relevant items can be identified. There is support for variable relationships between them.
E6	Formal description of architecture and security enforcing functions with correspondence between formal specification through source code and tests. All TOE configurations defined in terms of the architecture design and all tools can be controlled.

COMMON CRITERIA

Describe Common Criteria.

What do you get when you buy a CC (Common Criteria) evaluated product? These products have been through a level of testing and confirmation of some of their security strengths. The level of the evaluation indicates the type of testing done, but you get no guarantee that this product is free from exploitable vulnerability.

Moreover, you must realize that any product is certified by version and by environment. That is, even if a product is certified, this may mean nothing to you. You need to ask yourself three questions:

- ◆ Which version is certified? Is it the one I am using (or purchasing)?
- ◆ Is the environment where this product will be used the same as the one in which it was evaluated?
- ◆ Are the things this system was tested for important to my needs? Are there things not addressed by the certification?

If the first two questions are true, and the final one is satisfactory, you must still remember that the successful evaluation is only a measure of the extent to which security has been assessed. Keep this in mind as you study the Common Criteria.

What Is Common Criteria?

The “Arrangement on the Recognition of Common Criteria Certificates in the Field of IT Security” was signed as a mutual recognition arrangement in 1998 by government organizations from the United States, Canada, France, Germany, and the United Kingdom. This international standard, commonly known as Common Criteria, has as its objectives:

- ◆ Ensure IT product evaluations are performed to high and consistent standards
- ◆ Guarantee that evaluations contribute to the confidence in the security of the products
- ◆ Increase the availability of evaluated, security-enhanced IT products
- ◆ Eliminate duplicate evaluation
- ◆ Continuously improve efficiency and cost-effectiveness of security evaluations and certification/validation process for IT products and protection profiles

In sum, the CC provides an internationally agreed upon standard and evaluation methodology that can be used to certify IT products.

NOTE

Keeping Current The Version of the Common Criteria reviewed here is version 2.1, a version produced to align it to ISO/IEC 15408:1999, which can be downloaded from <http://csrc.nist.gov/cc/ccv20/ccv21list.htm> or <http://www.commoncriteria.org/cc/cc.html>. Additional associated modules, items that deal with areas not covered in the initial evaluation such as how to deal with flaws discovered in certified products, are also available from the Web site.

NOTE

Windows 2000 Windows NT, on which Windows 2000 is based, holds successful evaluations at the U.S. Orange Book C2 and the UK FC2/E3 IT Security Evaluation. Windows 2000 has been submitted for Common Criteria evaluation (EAL 4) for network operating systems. This includes evaluation of Windows 2000 Professional, Windows 2000 Server and Advanced Server, domain controllers, and advanced functionality (domain-based policy management, directory services, IPSec Services, Encrypting File System (EFS), and recovery services). Evaluation will be overseen and approved by the U.S. National Information Assurance Partnership (NIAP), the Common Criteria evaluation authority for the United States.

If, for example, a product passes an evaluation against the Common Criteria in England, it does not need to be tested in the United States.

The Common Criteria is divided into three parts:

- ◆ **Part 1: Introduction and General Model**—General concepts, principles of IT security evaluation, high-level specification writing, usefulness for target audiences. Good background and reference for consumers.
- ◆ **Part 2: Security Functional Requirements**—Functional requirements, components, Targets of Evaluation (TOEs); good for guidance and references consumers can use to formulate requirements for security functions.
- ◆ **Part 3: Security Assurance**—Assurance requirement for TOE's and evaluation criteria for Protection Profiles and Security Targets. Guides consumers on required levels of assurance.

Within sections two (Security Functional Requirements) and three (Security Assurance), a number of classes are defined. Classes are a general grouping of similar security functional requirements. Each class may be divided into one or more families or subdivisions. A family is a collection of requirements that share objectives, but each has a different emphasis or strength. Assurance families, however, have hierarchical components, while in Security Functional Requirements, classes may be hierarchical.

Each part of the Common Criteria is discussed in more detail in the following sections.

Part 1: Introduction and General Model

Part one provides definitions; thoughts on how the CC can be used by consumers, developers, evaluations, and others; the general model of the CC; and the requirements of the CC. Two important parts of any CC submission are the definition of a Security Target (this is the specification against which the product will be evaluated) and the Protection Profile (the security profile that the security target seeks to address).

The Protection Profile (PP) describes security requirements and indicates the security problem that the TOE will solve. Within the PP, CC functional and assurance requirements are stated along with a rationale for these components. An EAL (evaluation) may also be stated. Evaluation of PP can also be sought separate from the product evaluation. Criteria is stated in part 3. A PP evaluation indicates that the PP can be used as a statement of requirements for an available TOE.

A PP can be as simple as a company describing a security requirement for its e-commerce site, or as complex as a proposal to allow a presidential election to proceed on the Internet. They can also set a standard for a particular product type, such as a firewall. Many PPs have already been written, and a number of approved PPs can be located at the following:

- ◆ The Protection Profile PP registry www.radium.msc.mil/tpcp/library/protection_profiles/index.html.
- ◆ www.cesg.gov.uk/chtml/ippr/list_by_type.html
- ◆ csrc.nist.gov/cc/pp/pplist.htm
- ◆ Links page at csrc.nist.gov/cc/linklist.htm
- ◆ Scheme (country specific implementation body) sites

A Security Target (ST) is the basis against which evaluation is done. It contains the TOE security threats, objectives, requirements, and a summary specification of security functions, assurance functions, and assurance measurers. Another use for the ST is that a consumer can see whether security functionality of a product and its assurance package meet his requirements, and if its stated configuration is consistent with his proposed environment. ST evaluation criteria is also specified in part 3. Evaluation indicates its suitability for use as basis of its corresponding TOE. If it claims correspondence to a PP, evaluation demonstrates that it meets these requirements.

Part 2: Security Functional Requirements

Security requirements for a trusted product or system can be developed by considering the threat to IT. The components of the CC can be catalogued to create a security requirements definition.

The components are represented by eleven functional classes each of which is divided into families. The Security functional requirements are used to create the functional requirements of the TOE:

- ◆ **Audit (FAU)**—Security events are recognized, recorded, and analyzed to produce audit records. These records can be examined to determine security relevance. The audit class is divided into families. Each family defines what is an auditable event, and how records are analyzed, protected, and stored.
- ◆ **Cryptographic Support (FCS)**—Two families, one for operational use and the other for management of cryptographic keys, make up this class.
- ◆ **Communication (FCO)**—This class is concerned with assuring identity of parties involved in data exchange. One family is concerned with non-repudiation of the originator and the other of the receipt.
- ◆ **User Data Protection (FDP)**—The families within this class specify how user data is to be protected during import, export, and storage. Security attributes of data are also detailed.
- ◆ **Identification and Authentication (FIA)**—Identity of authorized users should be determined unambiguously. Security attributes associated with users and subjects need to be correctly associated. Families determine and verify user identity, their authority to interact with the target, and correct association of security attributed with users.
- ◆ **Security Management (FMT)**—Specifies management of security attributes, data, and function. Management roles (separation of capability) are defined. Covers management aspects of other function classes.
- ◆ **Privacy (FPR)**—Privacy requirements, including anonymity, anonymity with accountability, and so on. Protection of the user—preventing discovery and misuse of identity by other users.
- ◆ **Protection of the TSF (FPT)**—Protection of TOE Security Functions (TSF) data. Integrity and management, CIA, trusted recovery, replay detection, domain separation, time stamps, and so on.

- ◆ **Resource Utilization (FRU)**—Availability of resources: processing, storage capacity. Details for fault tolerance, service priority, resource allocation.
- ◆ **TOE Access (FTA)**—Control establishment of user's session. Limit number and scope of session, displaying access history, modification of access parameters.
- ◆ **Trusted Path / Channels (FTP)**—Trusted communication paths between users and TSF, and between TSF and TSF. Trusted channels exist for this purpose. An exchange can be initiated by user or TSF and is guaranteed protected from modification by untrusted applications.

Part 3: Security Assurance Requirements

Assurance, the demonstration that proposed security measures are sufficient to fulfill an organization's security policy and clearly articulated security threats, is defined for PPs, STs, and TOE. Two classes describe assurance requirements for PP (APE) and ST (ASE) evaluations whereas seven describe evaluation assurance requirements. One class describes assurance maintenance. The classes are

- ◆ **Protection Profile Evaluation (APE)**—Demonstrates that the PP is complete, consistent, and technically sound and states the requirements for an evaluable TOE. This should include information on TOE Description, Security environment, security objectives and TOE security requirements.
- ◆ **Security Target Evaluation (ASE)**—Demonstrates that the ST is complete, consistent and technically sound. It is suitable for TOE evaluation. This should include TOE description, security environment, PP claims, TOE security Requirements and TOE summary Specification.
- ◆ **Configuration Management (ACM)**—Integrity of TOE is preserved, TOE and documentation used for evaluation that is distributed.
- ◆ **Delivery and Operation (ADO)**—Security protection of TOE is not compromised during delivery, installation, and operations use.

- ◆ **Development (ADV)**—Refinement of TSF from ST specification to implementation. A mapping from security requirements to a low-level representation.
- ◆ **Guidance Documents (AGD)**—Secure operations use of TOE by admins and users.
- ◆ **Life Cycle Support (ALC)**—This class includes the lifecycle definition, tools, techniques, security of development environment, and the correction of flaws found by consumers.
- ◆ **Tests (ATE)**—TOE meets the functional requirements in this class. Examines the depth of developer testing as well as independent testing.
- ◆ **Vulnerability Assessment (AVA)**—Identification of exploitable vulnerabilities introduced by construction, operation, misuse, or incorrect configuration. Uses covert channel analysis, analysis of configuration, strength of mechanisms of security function, identifies flaws.
- ◆ **Maintenance of Assurance (AMA)**—Requirements the product should meet after certification as measured against the CC. Need to assure the TOE will continue to meet security target as changes are made to it or its environment. This provides a way to establish assurance maintenance schemes.

Evaluation Assurance Packages or Levels

EALs are combinations of assurance components. They also can be conveniently compared to TSCEC and ITSEC. Like these security evaluation criteria, EALs are scaled with from EAL1 through EAL7. Other EALs exist, but EAL7 is the highest with international recognition:

- ◆ **EAL1**—Functionally tested—Confidence in correct operation is required but threats are not serious. Due care has been exercised with respect to protection.
- ◆ **EAL2**—Structurally tested—Delivery of design information and test results are consistent with good commercial practice. Low to moderate level of independently assured security. Many legacy systems can be evaluated at this level.

- ◆ **EAL3**—Methodically tested and checked—Security engineering at design states, requires minimal alteration of existing sound development practices to meet. (Grey box testing, search for obvious vulnerabilities.)
- ◆ **EAL4**—Methodically designed, tested, and reviewed—Use of positive security engineering, good commercial development practices, rigorous, but does not require substantial specialist knowledge, skills, or testing. Independent search made for obvious vulnerabilities.
- ◆ **EAL5**—Semi-formally designed and tested—Semi-formally tested using rigorous commercial development practices, application of specialized security engineering techniques. High level independently assured security in planned development, rigorous developmental approach.
- ◆ **EAL6**—Semi-formally verified, designed and tested—Specialized security engineering techniques in rigorous development environment. Protection of high value assets against significant risks. Modular, layered approach to design, structured presentation of the implementation. Independent search for vulnerabilities ensures resistance to penetration, systematic search for covert channels, development environment, and configuration management controls.
- ◆ **EAL7**—Formally verified, designed, and tested—This is used for extremely high risk situations, or high value of assets. White box testing is used.

Areas Not Addressed by the Common Criteria

CC does not test secure usage. No assumptions are made about administration unrelated to IT security awareness in the organization. There is no evaluation of organizational, personnel, physical, or procedural controls. The following list specifies areas that that CC does not cover:

- ◆ Electromagnetic control is not addressed.
- ◆ Procedures for accreditation (this is an administrative process).
- ◆ Criteria for assessment of cryptographic algorithms not covered.

REVIEW BREAK**A Comparison of the Orange Book, ITSEC, and Common Criteria**

Table 6.6 lists the various classes or levels of Orange Book, ITSEC, and CC in a way that allows easy comparison. This model should serve as a reference to help those familiar with earlier evaluation or certification criteria. It does not mean that a one-to-one correspondence exists between every stitch at each level. It is more useful as an aid for those security professionals who are getting started with CC, rather than as a direct comparison tool.

TABLE 6.6
STANDARDS COMPARISON

<i>Orange Book</i>	<i>TCSEC</i>	<i>Common Criteria Evaluation</i>	
		<i>ITSEC</i>	<i>Assurance Level</i>
D	Minimal Protection	E0	EAL0 EAL1
C1	Discretionary Security Protection (discretionary access control, identification and authentication, system architecture, system integrity, security testing, documentation)	F1+E1	EAL2
C2	Controlled Access Protection (object reuse, and audit)	F2+E2	EAL3
B1	Labeled Security Protection (labeling, label integrity, design verification)	F3+E3	EAL4
B2	Structured Protection (covert channel, device labels, subject sensitivity labels, trusted path, trusted facility management, configuration management)	F4+E4	EAL5
B3	Security Domains (intrusion detection; security administrator role definition)	F5+E5	EAL6
A1	Verified Design (verified design, more documented version of B, trusted distribution)	F6+E6	EAL7

IPSec**Describe the Internet Protocol Security (IPSec) standard.**

The Internet Protocol Security standard (IPSec) is an IETF standard that describes a communications protocol that can be implemented.

TCP/IP, the original protocol developed for the Internet, and now the primary protocol for internal communications within IT networks and across WAN links, was not designed with security in mind. The protocol was developed with the goals of guaranteed connectivity and availability.

IP Security was originally designed for the future implementation of the Internet Protocol, IPv6, but is now specified for the current version of IP, IPv4, as well. Numerous implementations exist including built-in and add-on functionality for routers and firewalls, as well as packages for client computers. In addition, all versions of Windows 2000, Windows XP Professional and Windows .NET operating systems have built in IPSec capability.

Uses for IPSec

The primary uses of IPSec today involve the implementation of a Virtual Private Network (VPN) or the protection of communications between two computers, or a computer and a security device on the same LAN. However, IPSec can also be used to allow or block specific computers or communications protocols from entering or leaving a computer. An additional use can be authentication only. By using the AH protocol and requiring certificates for authentication, an administrator can control which machines can communicate with others on a network.

When used for communications between computers, either tunnel mode (VPN) or transport mode (communications between two computers on a LAN) IPSec provides the following:

- ◆ **Access control**—Access can be restricted by identifying the IP address of the computer(s).
- ◆ **Connectionless integrity**—A checksum is calculated and a hash is computed across the payload and is also encrypted.
- ◆ **Mutual computer authentication**—Prior to data transmission, each computer must authenticate to the other. The standard allows a multiple authentication technique. Implemented products use certificates, shared keys, or Kerberos.
- ◆ **Confidentiality**—The information is protected during transit. If the information is captured, it cannot be easily interpreted as it is encrypted.

- ◆ **Data-origin authentication**—Each packet can be attributed to the sending computer.
- ◆ **Protection against replay attacks**—Three items identify the communication, a Security Parameters Index (SPI), which identifies the appropriate Security Association (connection), a sequence number, and the authenticated computer's IP address. This information is kept on received data and no triplet should match one that has already been recorded. If it does, IPSec considers this an attack and drops the packet.

When used in *allow* or *block* mode, IPSec is configured on a single computer to narrow the types of communication that are acceptable. In normal TCP/IP communications, all data directed specifically to the computer, as well as data which is part of a broadcast (directed to all computers that are listening) is accepted by the network card and passed up the TCP/IP stack. The TCP/IP stack is a collection of protocols which add or remove communication-related information to the raw data sent from computer to computer. (A detailed discussion on how this works can be found in Domain 3, or Chapter 2 of this book, “Security.”)

IPSec can be configured to block specific protocols, such as HTTP (Web), FTP (file transfer), SMTP (mail), and so on, from either leaving or entering the stack. Let's say, for example, that Alice has installed a Web server on her Windows 2000 Professional system. This happens to be strictly prohibited by corporate policy, but Alice has done so regardless of the policy. This Web server now becomes a good target for Code Red and other Web server-related attacks. A blocking policy that blocks receipt of traffic to port 80 (the standard port for HTTP traffic) will prevent access to the Web server. It will not block Alice's access to the Internet, as traffic from her machine to port 80 on a Web server is not blocked.

Architectural Components of IPSec

IPSec is a modular protocol that uses Internet Key Exchange (IKE) for master key creation. The master key is used to create session keys, the keys used for encryption.

IPSec is composed of two subprotocols—IP Authentication Header (AH) and Encapsulation Security Payload (ESP). Both subprotocols can provide integrity, data origination authentication, mutual computer authentication, and anti-replay. ESP can also provide confidentiality.

Two phases are used to set up IPSec sessions. In phase I, after machine authentication, a security association (SA) is created and used for the exchange of keying material. IKE is used and a master key is created. The master key is used in phase II, and it can be required that the master key be recalculated periodically.

In phase II, session keys are created and two SAs are established. One is used for data traveling from computer 1 to computer 2, and the other for data traveling in the opposite direction. Session keys can be set to be renewed at periodic intervals. By frequently changing session keys (and if required, changing the master key), the risk of compromise is reduced. An attacker who deduces an encryption key will only be able to use it to decrypt captured information until the key is changed; then he must start all over again. Frequent rekeying can cause performance problems.

CASE STUDY: C2 AND WINDOWS NT

ESSENCE OF THE CASE

Three issues are at work here:

- ▶ First, a security evaluation should match the intended use of the product.
- ▶ Second, administrators should not run tools without understanding what they will do to systems. Adequate documentation explained the C2 certification for Windows NT 3.51 and should have alerted all but the clueless administrator as to what might happen when the tool is applied.

SCENARIO

Windows NT 3.51 was evaluated at the C2 level some years ago, and yet its networking components were not evaluated. A C2 tool was included with the Windows NT Resource Kit. Administrators wanting to have secured systems used the tool. When run, the tool configured the Windows NT system to meet the C2 requirements as established in the evaluation. I think you can guess what happened. When the unwary administrator used the tool, networking components were removed and the system became compliant with the specifications as evaluated, but was useless on the network. Moreover, since the administrator made no effort to understand exactly what the tool was doing, his troubleshooting efforts were compounded.

continues

CASE STUDY: C2 AND WINDOWS NT

continued

- ▶ Finally, in this tool-crazy point-and-click administration world, perhaps the release of such a tool was a little premature. Yes, admins should understand what they are doing, but tools can also adequately prompt users and provide a way to remove the effects of the tool's application.

ANALYSIS

Fortunately, Windows NT 4.0 was evaluated with networking components. In fact, six different configurations of NT 4.0 were evaluated. A handy administrator's guide to what the evaluation means, and how to implement C2 level security on NT 4.0 is available. The guide is thorough, and provides information on C2, the Orange Book and the specifics on the evaluation process and results. I especially like the emphasis on understanding the difference between configuring a system to C2 level, and obtaining accreditation as a C2 level site, in which they state:

“Please keep in mind that there is a difference between deploying a system in a C2-evaluated configuration and having a C2-certified system. A C2 evaluation considers whether a particular product (in this case, Windows NT) can be part of a C2 certification, when configured appropriately. A C2 certification indicates the degree of security that an actual deployment provides, and considers physical security, administrative procedures and other factors in addition to how Windows NT is configured. There can be considerable value in deploying Windows NT in one of the evaluated configurations, not the least of which is that doing so makes it eligible for certification. However, only an accredited certification facility can grant certification.”

I recommend you obtain a copy of the guide even if you have no intention of configuring NT 4.0 to the C2 level. You will learn much about the evaluation process and how it can be used. You can download the guide from <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/prodtech/cyrpto/c2deploy.asp>.

CHAPTER SUMMARY

This chapter covered various security architecture topics, from the design of access control, to the evaluation of computing systems by international standards. Additionally, a communications protocol standard was introduced from the perspective of security architecture suitable for network communications. Understanding security architecture may just be the lynch pin of future secure computing efforts. Because it is impossible to anticipate future attack substance, those responsible for IT security must return to the philosophy of securing systems first with known security practices, and then later in response to attacks which cannot be met by them.

KEY TERMS

- Assurance
- Bell-LaPadula model
- Biba Model
- Channel
- Clark-Wilson model
- Clearance
- Closed system
- Common Criteria
- Compartmentalization
- Covert channel
- Covert storage channel
- Covert timing channel
- Discretionary access control (DAC)
- Discretionary security protection
- Evaluation assurance level (EAL)
- Execution domain
- Export of labeled information
- Formal security model
- Formal verification
- Information label
- ITSEC
- Labeling
- Layering
 - Abstraction
 - Process isolation
 - Least privilege
 - Resource access control

CHAPTER SUMMARY

- Mandatory access control (MAC)
- Multilevel device
- Open system
- Protection profiles
- Reference monitor
- Secure state
- Security function
- Security kernel
- Security level
- Security model
- Security target
- Segmentation
- Sensitive information
- Single level device
- Star property (* property) or confinement property
- Target of evaluation (TOE)
- Trusted channel
- Trusted computing base (TCB)
- Trusted distribution
- Trusted facility management
- Trusted facility manual (TFM)
- Trusted path
- Trusted system
- TSEC
- Validation
- Verification

APPLY YOUR KNOWLEDGE

Exercises

6.1 Real-World Security Architecture Evaluation

Estimated Time: 30 minutes

1. Find a real-world example of a security architecture system that is in place and describe it.
2. Compare your assessment to the Windows 2000 assessment provided here.

The Windows 2000 security architecture is composed of the following components:

- ◆ **Local Security Authority**—A protected subsystem which maintains information about local security for a system. It also provides translation between names and identifiers, provides interactive authentication services, generates access tokens, and manages the audit policy and settings.
- ◆ **Net Logon Service**—Passes user credentials to the domain controller through a secure channel and returns domain SIDs and user rights. Maintains that channel.
- ◆ **Security Accounts Manager Service**—The service that enforces local security policies.
- ◆ **Security Reference Monitor**—Arbitrates access control on the system. User credentials must match the access control lists assigned to resources in order to use them.
- ◆ **Authentication protocols**—Kerberos v5 authentication protocol and NT LAN Manager (NTLM) authentication protocol, Secure Sockets Layer (SSL) authentication protocol.

All components of the security subsystem track security policies and accounts in use on the system. Accounts in a domain are stored in the Active Directory, whereas local system accounts are stored in the Security Accounts Manager (SAM).

Review Questions

1. How have past differences in public versus government requirements for security architecture affected evaluation criteria and security models that are in use today?
2. Compare the management of integrity by the Biba and Clark-Wilson models.
3. What is a reference monitor and why is it important?
4. Describe the difference between an open and a closed system.
5. Explain domain separation.
6. List possible uses for IPSec.
7. What does the common criteria not address?

Exam Questions

1. You apply the specified C2 level configurations to all of your Windows NT 4.0 Workstation and Servers. You now have
 - A. A C2 level accredited site
 - B. A C2 level certified site
 - C. A C2 level configured computer
 - D. No networking functionality

APPLY YOUR KNOWLEDGE

2. Which security model addresses only confidentiality?
 - A. Bell-LaPadula
 - B. Biba
 - C. Clark-Wilson
 - D. Access control lists
3. A flaw that allows an object with legitimate access to information to transfer the information in a manner that violates system security policy is a what?
 - A. Limited access mode
 - B. Backdoor
 - C. Multi-level access system
 - D. Covert channel
4. The certification emphasis of TCSEC is
 - A. Confidentiality
 - B. Availability
 - C. Integrity
 - D. Least Privilege
5. The B level of TCSEC certification is important because it is at this level that the concept of _____ is introduced.
 - A. Auditing
 - B. Accountability
 - C. Labels
 - D. Separation of users and data
6. The C2 level of TCSEC certification is important because it is at this level that the requirement for _____ is introduced.
 - A. Auditing
 - B. Labels
 - C. Covert channels
 - D. Configuration management control
7. The ITSEC security architecture addresses what?
 - A. Confidentiality
 - B. Assurance, confidentiality, integrity, and availability
 - C. Confidentiality, integrity, and availability
 - D. Integrity
8. A document that describes security requirements and indicates the security problem that the TOE will solve is the what?
 - A. Security target
 - B. The protection profile
 - C. A security functional requirement
 - D. Covert channel
9. IPSec is composed of which two subprotocols?
 - A. AH and ESP
 - B. TCP and IP
 - C. FTP and TCP
 - D. FTP and IP

Answers to Review Questions

1. Government requirements for security have been centered around confidentiality. Therefore, many of the early security models (Bell-LaPadula) and architecture (TCSEC) have had confidentiality as their major emphasis. See the “Requirements for Security Architecture and Models” section for more information.

APPLY YOUR KNOWLEDGE

2. The Clark-Wilson Model is directed towards commercial enterprise versus the government focus of Biba. Thus Biba focuses on a lattice of integrity and a no write up and no read down model whereas Clark-Wilson focuses on software engineering concepts such as abstract data types, separation of privilege, allocation of least privileges, and non-discretionary access controls. It also addresses the issues of authorized users making modifications they are not authorized to do, and preventing unauthorized user from making modifications. See the “Security Models” section for more information.
3. A reference monitor is an abstract concept that stands for the arbitration of access to resources. It is important because it is one of the requirements of secure systems. See the “Reference Monitor” section for more information.
4. An open system gives all users administrative level access. It also uses standard user interfaces. A closed system is totally secure. It does not use standard user interfaces. See the “Open Versus Closed Systems” section for more information.
5. Domain separation is a function of the system design. It means that functions are grouped according to their purpose and need to access each other and defined resources. This grouping is called a domain of that function. Domain access is restricted. For example the user functions do not need to access the kernel, so they are not allowed to access that domain. See the “Security Architecture” section for more information.
6. IPSec can be used for confidentiality, data origin authentication, protection against replay, mutual authentication, integrity, and access control. See the section “Uses for IPSec” for more information.
7. Common Criteria does not address issues of electromagnetic control, procedures for accreditation, or for assessment of cryptographic algorithms. See the section “Areas Not Addressed by the Common Criteria” for more information.

Answers to Exam Questions

1. **C.** Applying the steps to configure an evaluated computer to the level at which it has been evaluated only does that—configure it to the same level it was evaluated at. To have an accredited site, you must obtain accreditation from an accreditation body. See the case study and the “TSCEC: The Orange Book and the Rainbow Series” section for more information.
2. **A.** The Bell-LaPadula security model only addresses issues of confidentiality. See the “Bell-LaPadula” section for more information.
3. **D.** A back door is a planned access channel to a system. Multi-level and limited access modes are specific operational modes that a system may have. See the “Covert Channel” section for more information.
4. **A.** TCSEC does not address availability, integrity or least privilege. See section “TCSEC: the Orange Book and the Rainbow Series” for more information.
5. **C.** The other three are introduced in level B. See the “Orange Book Classifications” section for more information.
6. **A.** The others are addressed at level 2. See the TCSEC: The Orange Book and the Rainbow Series” section for more information.

APPLY YOUR KNOWLEDGE

7. **B.** All these issues are addressed in the standard. See the “Differences between the Orange Book and the ITSEC” section for more information.
8. **B.** The protection profile (PP) is the answer. The security target defines the evaluation criteria that should be met. Security functional requirements are the individual classes defined in section 2 and covert channels are defined earlier. See the “Part 1: Introduction and General Model” section for more information.
9. **A.** AH and ESP is the correct answer. TCP and IP are components in TCP/IP, FTP is the File Transfer Protocol. See the section “Architectural Components of IPSec” for more information.

Suggested Readings and Resources

1. Bragg, Roberta. “New Products, Protocols, and API.” In *Windows 2000 Security*. New Riders, 2000.
2. <http://csrc.nist.gov/cc/ccv20/ccv2list.htm> (Common Criteria).
3. <http://Csrc.nist.gov/cc/pp/pplist.htm> (protection profiles).
4. <http://www.cesg.gov.uk/assurance/iacs/itsec/index.htm> (ITSEC’s Web site).
5. http://www.cesg.gov.uk/chtml/ippr/list_by_type.html (protection profiles).
6. <http://www.commoncriteria.org/cc/cc.html> (Common Criteria).
7. <http://www.microsoft.com/technet/treeview/default.asp?url=/TechNet/security/prodtech/cyrpto/c2deploy.asp> (C2 and Windows NT).
8. http://www.microsoft.com/windows2000/techinfo/reskit/en-us/default.asp?url=/WINDOWS2000/techinfo/reskit/en-us/distrib/dsbg_dat_doq.asp (Window 2000 security architecture).
9. http://www.radium.mcsc.mil/tpep/library/protection_profiles/index.html (protection profiles).
10. <http://www.radium.ncsc.mil/tpep/library/rainbow/> (TCSEC).

OBJECTIVES

Identify the key roles of operations security.

- Identify resources to be protected.
 - Identify privileges to be restricted.
 - Identify available controls and their type.
 - Describe the OPSEC process.
- Loosely defined, operations represent the “do.” For a business, this can mean production, distribution, transportation, and any function that gets the main work of the organization done. Operations security then, could be considered the mechanisms that protect these essential functions. However, for many of us, operations security has come to mean the protection of the computing infrastructure, the information it processes and its input and output. I will pursue this meaning, but to do so I’ll first define its key roles.

Define threats and countermeasures.

- The first question to be asked is, “What are we securing our computer operations from?” It is impossible to secure anything from the vague uneasiness and requirement that security must be imposed. Listing perceived threats will focus the thought process and enable discussion of which threats are probable, or possible; which vulnerabilities are present in our systems; and how we are at risk. Next, we can select *countermeasures*, or ways to prevent or mitigate the risk that threats will succeed.



CHAPTER 7

Operations Security

OBJECTIVES

Explain how audit and monitoring can be used as operations security tools.

- **Explain how audit logs can be used to monitor activity and detect intrusions.**
 - **Discuss intrusion detection.**
 - **Explain penetration testing techniques.**
- ▶ Security is not something you do when it pleases you. It is not simply hardening systems (a process of removing unnecessary elements and configuring others to make the system as secure as possible), applying patching, and configuring a firewall. Security is a continual process. One part of the process is monitoring for abnormal events, unapproved changes, and other potential symptoms of attack. Three primary methods of monitoring are audit (are things as they should be?), intrusion detection (is somebody attempting to steal or change things?), and penetration testing (can a friendly force get past your security?).

Define the role of Administrative management in operations security.

- ▶ Security officers are not the only ones who should be involved in keeping operations secure; each employee has a role to play. Management, however, has a special, key part to perform. Management must be the lynchpin, the element that both connects the activities of others and holds the parts in place. Three management roles that impact security are policy, employee supervision, and expenditure approval.

Define operations security concepts and describe operations security best practices.

- **Explain antivirus controls and provisions for secure email.**
 - **Explain the purpose of data backup.**
 - **Detail how sensitive information should be handled.**
 - **Describe how media should be handled.**
- ▶ Although all operations should be scrutinized to determine the policies and procedures that will best keep them secure, several concepts are so key to survival that they need specific mention:
- Email has become a mission critical operation; every procedure possible should be used to keep it secure.
 - Backup remains the one consistent recovery strategy. Without a solid backup plan, every organization's data is at the mercy of a hardware glitch or environmental disaster.
 - Sensitive information requires special handling, but does everyone agree on what information is sensitive?
 - Media (tapes, discs) is not indestructible. How can you ensure that the media you use will keep your data safe?

OUTLINE

Introduction	385	Establishing Countermeasures for Employee-Related Threats	412
Examining the Key Roles of Operations Security	387	Including Countermeasures in Hiring and Firing/Exit Practices	414
Identify Resources to Be Protected	387	Gruntling Program	415
Identifying Privileges to Be Restricted	388	Countermeasures for Common Internet-Based Threats	416
Identifying Available Controls and Their Types	389	Countermeasures to Physical Threats	417
Control Types	391		
Describing the OPSEC Process	391	The Role of Administrative Management	418
The Roles of Auditing and Monitoring	395	Concepts and Best Practices	420
Using Logs to Audit Activity and Detect Intrusion	396	Privileged Operation Functions	421
Detecting Intrusions	399	Understanding Antiviral Controls	423
Penetration Testing Techniques	403	Protecting Sensitive Information and Media	425
Developing Countermeasures to Threats	408	Change Management Control	427
Risk Analysis	408	Chapter Summary	430
Threats	409	Apply Your Knowledge	432
Countermeasures	411		

STUDY STRATEGIES

- ▶ Operations security covers a lot of ground. From management of equipment to management of people, the topics it relates to have no end. One of the challenges of understanding this broad topic is its reliance on the underlying technology. Concepts and best practices will not make sense if you do not understand how computers, networks, programs, data centers, and businesses work. If you do not already have some experience with them, try to find someone who does and ask them to help you understand. Spending some time with Chapter 2, “Telecommunications and Network Security,” will help as well.
- ▶ If your background is not in technology, mastering the material in the domains of Security Management, Telecommunications and Networking Security, Disaster Recovery and Business Continuity, and Application and Systems Development is essential. Study them first, and this chapter will be easier. If your background is technical, do not be frustrated by the light treatment of technical content here. Operations security is more concerned with the big picture than with the intimated details of how to configure or code.
- ▶ Whatever your current job description, whatever your background, use your knowledge of this domain to do two things:
 - First, see what you can determine about operations security at your organization. If you aren’t working in a directly related area, don’t be surprise if you can’t find out much. Good operations security is transparent; that is, it reveals little about the specifics of its activities.
 - Second, operations security principles can be applied to things other than computer operations. Military organizations have long used these principles to improve their prospects of success. Practice these principles on your activities on the Internet. What can someone learn about you while you’re online? How might that be used to their advantage? What can you do to diminish the amount of information that can be gleaned from your activities?

“Operations security is used to identify the controls over hardware, media, and the operators with access privileges to any of these resources. Audit and monitoring is the mechanisms, tools and facilities that permit the identification of security events and subsequent actions to identify the key elements and report the pertinent information to the appropriate individual, group or process.

The candidate will be expected to know the resources that must be protected, the privileges that must be restricted, the control mechanisms available, the potential for abuse of access, the appropriate controls, and the principals of good practice.

—Common Body of Knowledge study guide

This chapter covers Domain 7, Operations Security, which is one of 10 domains of the Common Body of Knowledge (CBK) covered in the Certified Information Systems Security Professional Examination. This domain is divided into several objectives for study.

INTRODUCTION

Operations security is the combination of two practices. It is the implementation of sound security principles and the gleeful application of a paranoid viewpoint to day-to-day operations.

There are many lists and papers that discuss the how and why of hardening systems and securing data. We know, in general, the steps we need to take to secure our data, our systems and our network. We can provide reams of documentation that detail how to best handle tapes; keep dirt and dust out of the data center; avoid conflict of interest; reduce opportunities for fraud, embezzlement, and espionage; and secure OSes, applications, and hardware. This general security takes us a long way. But it is the second practice, the activity which stems from seeing through the eyes of the enemy and operating as if “everyone’s out to get me” that moves security beyond the static application of practice to the daily strengthening of defenses.

Imagine you live in ancient times. Imagine you are a king out to conquer the world. You are not content to wait for your enemies to attack, but you are not willing to advance without knowledge of the enemies' strengths, weaknesses, and plans. Daily, you send spies to reconnoiter the territory, and daily you torture the captured to discover information about your enemy. You spend endless hours sketching his operations. How many horse soldiers does he have, and how many pikemen? Can his defenses stand up to your battering ram or catapult? Did your last attack deplete his forces, or are reinforcements close by?

One of your spies returns with the details of your enemies' defenses. Buckets of boiling oil await your next attempt at climbing the outer walls. The drawbridge is up and archers man the slits in the wall. Beneath the deep, dark waters of the moat lurk strange animals that occasionally break the surface with a fin or scaled side.

Suddenly, your reverie halts. Your body stiffens, recognition dawns. Your enemy must be studying you, as you study him. What actions might your troops be performing that give away your defenses, intentions, and vulnerabilities? Subtly, quietly, you change your modus operandi to mask what you are about.

In this way, operations security, or OPSEC, was born. *OPSEC* is the practice of looking at your sensitive operations through the eyes of your enemy and developing your security practices so he sees nothing. This natural complement to defensive measures has long been a practice of the military. Because you are concerned with computer security I will describe this process by looking at computer operations. However, OPSEC could be applied to all business operations, whether or not they include the use of computers.

To protect both data and computer operations, operations security must be strong in both general security and the practice of OPSEC. Organizations must follow the military practice: Develop a strong defense based on current knowledge, and improve those defenses by scrutinizing operations from the perspective of the enemy. This chapter explains how to achieve both approaches.

EXAMINING THE KEY ROLES OF OPERATIONS SECURITY

Identify the key role of operations security.

- **Identify resources to be protected.**
- **Identify privileges to be restricted.**
- **Identify available controls and their type.**
- **Describe the OPSEC process.**

The first step in fulfilling the promise of operation security is to understand its key role. Operations security starts by identifying

- ◆ Resources to protect
- ◆ Privileges to be restricted
- ◆ Controls necessary to do so

Identify Resources to Be Protected

The first step in any security review is to determine what you need to protect. Some common items in every organization's information systems should be protected, including

- ◆ Computers, including servers, desktops, and laptops
- ◆ Routers, switches, and other networking appliances
- ◆ Printers
- ◆ Databases, including the database management software and content
- ◆ Security software and appliances (firewalls, intrusion detection systems [IDSs], biometric devices, Public Key Infrastructure [PKI])
- ◆ Media such as tapes, CD-ROMs, and disks
- ◆ Personal digital appliances (PDAs), phones, and wireless devices
- ◆ Modems and other communications devices

- ◆ Software, including licensed commercial software and custom applications
- ◆ Source code
- ◆ Documentation
- ◆ People

Identifying Privileges to Be Restricted

For each asset listed, what can be done to protect it, and how is the item used? Think first in broad areas, such as simple use, installation, configuration, modification, granting of access to others, and full control. Describing *permission sets* (who can read, write, or execute the files, for example) for the different types of objects within your infrastructure is also necessary. Also, *privileges* (such as who can log onto the computer or who has the right to access it over the network) can be unique for the object type and perhaps even for the brand and version of the product.

Operating systems, for example, have many privileges associated with their use and management. Although collections of privileges can be given automatically to administrators (the root account in Unix, the Administrators group in Windows NT/2000), it is possible to assign individual privileges to a user or groups of users. Many of these privileges concern operational control, such as the right to logon or the right to shut down the system. Operating systems can also classify their own code and allow only certain instructions to access core, or privileged, areas of its own function. This area is often referred to as the *kernel* or *ring 0*, and the code used to access it as *privileged instruction*. User level code is not allowed to directly run instructions in this core area.

Data center operations are awash with the need to manage privileges. Who should enter them? What are privileges capable of when they are set? Who should back up the data? Who restores data? Who's responsible for monitoring the logs, configuring the firewall, and approving overtime? Each of these issues must be identified at the earliest stages of operation security.

Identifying Available Controls and Their Types

Controls are the means to prevent misuse or abuse of privileges while allowing authorized individuals or processes to do their jobs. When you require employees to enter a username and password, you are using a control to restrict access to your networks and by extension to the data on them. To make controls easier to discuss, they are commonly divided into types. Three different classification schemas are often used. One of these schemas shows control types listed as

- ◆ **Operational controls**—These are day to day procedures, mechanisms that include physical and environmental protection, privileged entry commands, change control management, hardware controls, and input and output controls.
- ◆ **Audit and variance detection controls**—These are audit logs that contain information on the exercise of privilege and/or records of system activity. Variance detection products detect and can send alerts when unusual activities occur. Intrusion detection systems fall into this category, as do special programs such as Jammer and Tripwire, both of which record changes to file systems and operating system configuration databases.
- ◆ **Application software maintenance controls**—These controls monitor installation and updates to applications, and they keep a record of changes.
- ◆ **Technical controls**—These controls audit and journal integrity validations, such as checksums, authentication, and file system permissions.
- ◆ **Administrative or management controls**—These control personnel screening, separation of duties, rotation of duties, and least privilege.

Input and output controls protect computers and applications by monitoring and rejecting or accepting data at these entrance and exit points. The infamous buffer overflow is a good example of the problems that occur when poor input controls exist. Although a more complete discussion can be found in the Application and System Development domain, it's important to note two things. First, a buffer overflow results when too much data is passed into a program or part of a program. It's the technology equivalent of having too much to eat or drink with the same unpredictable results.

NOTE

Pound. Pound. Pound. Do you remember the commercial advertising one telecommunications company's ability to bill for time periods of less than a minute? In the commercial, a gum-smacking grocery clerk weighed every vegetable or part of the vegetable as one pound and charged accordingly. I had a similar experience recently. I filled my shopping cart with vegetables and fruits. During check-out, I typically pack sacks instead of observing the clerks operations. This grocery had an automatic scale. The clerk has only to place the produce on the scale and punch in the unique code. The cash register does the math and adds it to your bill.

At one point the clerk remarked, "Boy, that's an expensive mushroom!" I had placed a single Portobello mushroom (\$5.99/lb) in my cart. It weighed about 1/2 lb. The price on the read-out was \$17.97. Long story short, the scale was broken and was weighing every fruit or vegetable as if it weighed three pounds! Management was apologetic and for my troubles gave me all my fruits and vegetables free. In short, the clerk served as an output control. She saw something out of the ordinary and called the entire operation into question.

Second, buffer overflows can be prevented by good coding practices. Output problems can be resolved by checks on the data against other known or plausible results. Checks and balances on accounting reports is an example; another might be that shipping tickets should be checked for accuracy and reasonableness. Should a shipping ticket be issued for 10 million cans of tomato soup when the normal customer shipment consists of a few cases—perhaps it should be questioned. Output, in other words, should be automatically compared to norms and Unusual results pulled for cross checks.

Another way of looking at controls is the set of controls listed here:

- ◆ **Deterrent controls**—These controls reduce the likelihood of attack.
- ◆ **Preventative controls**—These controls protect vulnerabilities, reduce the impact of attacks, or prevent an attack's success.
- ◆ **Detective controls**—Detective controls detect an attack and may activate corrective controls or preventative controls.
- ◆ **Corrective controls**—These controls reduce the impact of an attack.

Another methodology used is to describe the controls that are applicable to a particular piece of equipment or function. For example, if you wished to describe controls applicable to PCs you might list the following:

- ◆ Disk locks to prevent use of portable media such as floppy disks and CD-ROMs
- ◆ Training on how to use controls
- ◆ Required passwords for access (logon)
- ◆ Acceptable use policies including rules, such as prohibiting the illegal copying or installation of software
- ◆ Requiring virus checking on all disks before use (if policy allows using portable media)
- ◆ The use of antivirus software
- ◆ Checking for compliance
- ◆ Requiring file encryption

NOTE

Separation of Duties The separation of duties is a central concept to security. It means that no one individual has the ability to perform both halves of any task that would allow him to commit fraud or to steal money or information. The separation of duties is not just an information systems security principle; in fact, it is sometimes easiest to think of in terms of general activities that can be performed manually. For example, we wouldn't want the individual who has the ability to set security configuration on the system to also act as a systems programmer.

- ◆ Requiring biometrics for authentication
- ◆ Requiring that help desk or IT staff configure PCs, not users

R E V I E W B R E A K

Control Types

I'm sure you can come up with specific controls you might have implemented in the past, but can you then take each control and map it to the control types mentioned previously? Table 7.1 groups similar controls and identifies their type.

TABLE 7.1
CONTROL TYPES

<i>PC Control</i>	<i>Control Types from Different Schemas</i>	
Requiring passwords for access, requiring biometrics for authentication	Technical	Preventative
Disk locks	Technical	Preventative
Acceptable use policies, requiring virus check of portable media	Operational	Preventative
Checking for compliance	Audit and variance detection	Corrective
Using antiviral software	Technical	Preventative
Requiring file encryption	Technical	Preventative
Training in controls	Management	Preventative
Requiring that help desk or IT staff, not users, configure PCs	Management	Preventative
Software code audit looking for buffer overflows	Technical	Input, output
Loading a personal firewall/IDS system	Technical	Detective

Describing the OPSEC Process

“The whole point of operations security is to have a set of operational (daily, habit ingrained) practices that make it harder for another group to compile critical information.”

—http://www.nswc.navy.mil/ISSEC/Docs/Ref/GeneralInfo/opsec_basics.html

The OPSEC process is the process of understanding your day-to-day operations from a competitor's/enemies'/hacker's viewpoint and then developing and applying countermeasures. By studying the OPSEC principles, you can develop more effective defenses for your own systems. OPSEC applies five principles to do this:

- ◆ **Identify critical information**—This is information key to the survival of the troops (computer operations). In information systems, as in military operations, the people most familiar with the project can best determine the critical nature of its data. Two areas of concern are the operation of computer equipment and the processing of data by that equipment. The data owners can best assign the value of the data and the impact if lost, or obtained by competitors or perhaps by the public. Computer operations are more completely understood by those who are responsible for them. What information specific to the information system realm is sensitive and should be protected? What related information might not be sensitive, but might reveal important information that would assist an attacker? OPSEC practice calls this information *indicators*.
- ◆ **Analyze threats**—Next, determine what threats exist. A *threat* is the ability to do harm, coupled with the intention to do so. What nations would seek to attack yours? Do terrorists see targets in your very symbols of freedom and prosperity? Are people actively seeking the possession of information that you control? Are there people with the necessary skills to successfully attack your systems? Here you must identify your adversary and his capabilities and goals. An astute OPSEC person will allow the people who do the work, whether it be IT operations or departmental-level projects, to initially identify the threats. Then, the OPSEC-trained person performs the analysis.
- ◆ **Assess vulnerabilities**—Vulnerabilities are faults that can allow a threat to develop into a successful compromise and cause harm. How much information is publicly available? How is information stored, disseminated, manipulated, and destroyed? Are the systems that participate in these actions without fault? What faults are there?

- ◆ **Assess risks**—Could a vulnerability be manipulated by those that threaten your systems? Can the information be collected, processed, evaluated, analyzed, and interpreted in time to use it? What would the impact of its use be? By impact, I mean impairment of ability to offer normal services, destruction of facilities or some component therein, or chance of harm to individuals.
- ◆ **Apply countermeasures**—What are the solutions? Can the vulnerabilities be removed? Can the threats be mitigated? Can the risks be reduced? Specific solutions and their impacts should be detailed along with the expected reductions in risk that they offer. Countermeasures can be viewed as any action that removes or reduces information or access available to the enemy. These might be small changes to procedures, better control over information, increased traditional security, and deception. They can include the disruption of an adversary's ability to collect, process, and analyze this information.

OPSEC proponents tell us the importance of continually revisiting each of these principles. Even as you congratulate yourself on the application of countermeasures, new threats can be perceived, new vulnerabilities uncovered. This is why no list of hardening steps or penetration test will ever fully succeed in making a computer or network secure.

This is why OPSEC focuses on *indicators*, the information that can be seen, heard, or collected from Web sites, tapes, discs, and documents. Indicators can be simply observing how things are done or noting the hours of operation. They can even be the astute and familiar observer's notice of deviations from the norm. Knowledge of typical arrival and departure times for important staff, for example, allows the would-be attacker to deduce important occurrences when these normal patterns change. Knowledge of emblems, identification, acronyms used, number of visitors—all these items can provide the attacker with useful information. For example, observation of a large number of high-level officers (because I know their identifying rank indicators) arriving might mean eminent troop movements or new aggressive activity. Or, knowing the acronyms used by a software company for new products in design might allow me to quickly identify useful information in captured emails.

Tip-off indicators provide focus for the attacker by telling him where to concentrate his efforts. These might be an increased volume of visitors, increase in activity, arrival of important staff, the use of particular acronyms, and so forth. Tip-off indicators can also be technical in nature, such as the ability to determine the operating system of Web server type in use. For example, the ability to determine that a Web server is Microsoft IIS allows an attacker to select IIS-specific attacks. Another type of tip-off indicator might alert a potential attacker of your countermeasures to his attack, which would then allow him to develop countermeasures to your countermeasures. To learn more about indicators, check out this article at the Central Florida Industrial Security Awareness Council Web site: www.cfisac.org/resource/OPSEC%20Indicators.

IN THE FIELD

ISP REVEALS INTERNAL PROCESSING

I'm not going to provide you with the URL where I found this information, but I think its important to realize such information can readily be found on the Internet with little effort. Recently I was searching for something totally unrelated when I found a reference to some interesting ISP data. The keywords *password* and *policy* caught my attention. The link did not produce results, but the cached pages were still available from the search engine. Here's what I found:

- Descriptions of internal security measures
- How access for internal users could be obtained
- Copies of the forms used to request access to a customer mailbox and the procedure used to do so
- Names of members of the help team and what they had control over
- Who was responsible for maintaining account access
- Phone numbers for the security response center (this was advertised as the place for employees to go to reset passwords)

This, obviously, is not information that should be exposed on the Internet. Information like this could enable an attacker to impersonate an employee and possibly obtain access to confidential information. By the way, the information is now also removed from the cache.

In another unrelated incident, a casual search revealed the name of an individual conducting a review of operations security for a military group. The existence of a review of security is not in itself useful information. We know that they occur. However, this source also revealed the name of the person conducting the review, dates, the areas of his concentration, and the purpose, which was an analysis of future operations and threats with recommendations for improvement. The report would be a goldmine for an enemy, because it could be a future blueprint of security deployments. This tip-off indicator provides enough information on which to allow an attacker to focus attention.

Don't underestimate the power of the Internet for revealing information. *A good OPSEC technique is to constantly use Internet search engines on your company, its departments, and principles. You might be surprised at what you find.*

Some good countermeasures to possible risks of making indicators widely available are document shredding, not replying to unsolicited mail and requests for information, eliminating the indicators, camouflaging and concealing the activity, and preventing information viewing and destruction. Another, quite different technique is *counteranalysis*, confusing the enemy with misinformation.

THE ROLES OF AUDITING AND MONITORING

Explain how auditing and monitoring can be used as operations security tools.

- **Explain how audit logs can be used to monitor activity and detect intrusions.**
- **Discuss intrusion detection.**
- **Explain penetration testing techniques.**

Auditing is often defined as the process of checking current activity against policy. In the United States, a letter announcing that the Internal Revenue Service will audit you induces panic. Your entries on a tax return will be judged against a set of laws that few understand completely, that all must pretend to know to file their taxes, and that even experts disagree upon. An audit of your information systems compliance with security policy should be less stress inducing, at least where security policy is clearly defined.

Security configuration can be checked against the norm, and audit logs can be inspected for deviation. Audit logs are also useful to the systems and network administrators, who, as part of their daily review, can find evidence in them of potential attacks. Automated programs can also be trained to discover patterns that might indicate intrusion.

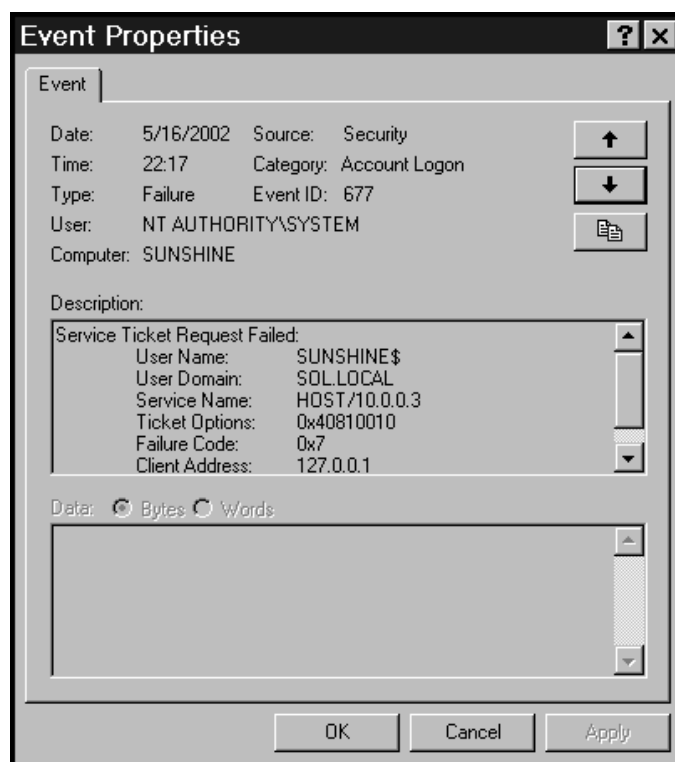
Following are the methods discussed in the next sections:

- ◆ Using logs to audit activity and detect intrusion
- ◆ Other methods of detecting intrusions
- ◆ Penetration testing techniques

Using Logs to Audit Activity and Detect Intrusion

Most computer systems are capable of logging information about operations occurring on them. In some cases, such as Windows NT/2000/XP, audit logging might have to be turned on and configured. The information that might be directly or indirectly found in the logs depends on the type of log and how it is configured. Logs can record operating system (information about the OS), application (information on applications running on this computer), and security (information on who is using the system and what they are doing) information. On some systems a single log can include all types of information; on others multiple logs of the same type, each one specific to an application, exist.

Figure 7.1 is a snapshot of a small portion of a security log on Windows 2000. Some of this information might be more useful to systems administrators because it records system operation and errors that can be used in troubleshooting. Other information is directly useful for auditing and intrusion detection. Logs can be analyzed to determine compliance with procedures, to provide a detailed audit trail of activity, to provide individual accountability, to enable the reconstruction of events, and potentially to detect an intrusion.

**FIGURE 7.1**

Audit logs present information related to security activity.

Table 7.2 lists the typical information found in logs, the type of log the information is found in, and how it might be used for auditing or intrusion detection purposes.

TABLE 7.2**WINDOWS 2000 LOGS**

<i>Information</i>	<i>Log Type</i>	<i>Discussion</i>
Record of system start up, normal shut down, and nonstandard shut down	System	Many attacks require physical access to the computer console and the ability to access maintenance modes or to boot to different OSes. Matching system shut down and start up to recorded maintenance events can reveal the existence of compromise attempts (or successes). Knowing that a system has been rebooted should trigger further investigation. In the mainframe world, start up is called <i>Initial Program Load (IPL)</i> . An unscheduled IPL might also be evidence of an attack.

continues

TABLE 7.2*continued***Windows 2000 Logs**

<i>Information</i>	<i>Log Type</i>	<i>Discussion</i>
Error messages about malfunction	System	Can be used in troubleshooting system problems.
Record of change in security policy	Security	Should be compared to manual authorized change logs to detect possible compromise.
Failure of application service to start	Application	Used in troubleshooting, this can be a symptom of something that might mean an attack and therefore should be monitored.
Successful logon	Security	Can be used in conjunction with logoffs and resource access to trace users' activity on the system. A successful logon after repeated logon failures might mean a successful attack.
Failed logon	Security	Can be symptomatic of a user who has forgotten her password, or it might record an attempt to break into an account.

Although one use of logs is to troubleshoot system performance or malfunction, the purpose of logs for operations security is to provide an audit trail and information that potentially points to an intrusion or breach in system security. Although the information in logs can't be reviewed in real time, it can provide evidence of attacks. Attackers often make several runs at a system rather than one large attempt; detecting an attack that occurred yesterday can be valuable in preventing the return attack that could come today.

Logging and log maintenance can present several problems. One might be lack of information, but many times the problem is that the information must be managed. Logs must be reviewed, and the entries must be evaluated—not just line by line, but with attention to the relationships between entries. For example, many logon failure records in the security log can indicate an attempt at cracking a user's password. Looking at a single line of the log might make the reviewer think that this is simply a case of a user making a typo when logging on. Many intrusion detection programs review security logs to detect an attack.

Detecting Intrusions

Intrusion detection is a technique used to identify intrusion attempts at and successful intrusions into a network or host machine. To understand intrusion detection techniques you must learn a little about how information from one computer travels to another. Just as voice communication over a telephone requires a conversion from words recognizable to the human ear to electric patterns that can flow across a wire, so data which can be viewed through application interfaces on a computing system must be modified for transmission. It must be changed into electronic signals that can travel between computer network interfaces and eventually reformed into data that can be used by the computer or by a human viewing the data through applications that reside on it.

Just as the destination computer can translate the data into a meaningful form to humans, the data traveling between computers can be captured and its patterns analyzed to determine its meaning. This can be useful to administrators who are troubleshooting, but it is also useful to attackers looking for information.

Various network monitors, intrusion detection devices, sniffers, and protocol filters can be purchased and run to collect or capture the data traversing the network. The data can then be analyzed.

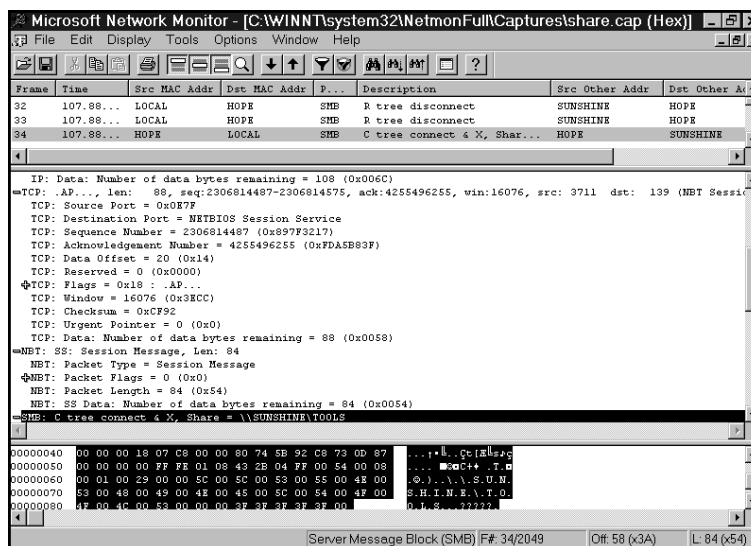
Hardware-based analyzers are attached to the network to listen to all communications. *Software-based sniffers* run on normal PCs. Because a PC normally only pays attention to data meant for it, the software-based sniffer alters the normal mode of the network interface card in the computer to “listen” to all the data on the network. The collected data is called a *capture*. This new mode of the network interface card is called *promiscuous mode*.

Figure 7.2 is a capture taken with Microsoft’s Network Monitor. The screen is divided into three sections. In the top section, all captured packets are listed. The highlighted packet is expanded in the middle section. Each + represents an area of information that can be expanded.

400 Part I EXAM PREPARATION

FIGURE 7.2

Information gathering using Network Monitor.

**NOTE**

Sniffing on Switched Networks You might hear that switched networks are more secure because it is impossible to sniff them. They're not. They are more difficult to sniff, but not impossible. The reason people think they cannot be sniffed is because switches, unlike hubs, deliver packets only to the systems they are addressed to. Thus, a sniffer placed on the network could not listen to all traffic, because not all traffic would be available on its portion of the network. However, techniques for sniffing switched networks have existed for several years. You can find an excellent introduction to the concepts of sniffing switched and nonswitched networks at http://www.sans.org/newlook/resources/IDFAQ/switched_network.htm.

In this view, the highlighted information clearly shows the name of a file share. On the second line, the source and destination ports are listed. The use of port 139, the Netbios Session Service, indicates that this connection is from a Windows computer. Port 445 indicates a session between two Windows 2000 computers that were members of the same forest, or at least between trusted domains, but the use of port 139 is only suggestive that this is not the case. It could be between a Windows NT computer and a Windows 2000 computer or some other combination of Windows computers. Much more information could be gained by further analysis of this packet and others in the capture.

The final segment of the screen shows the raw data capture. Even here, note the clear text file share name. If I were to spend more time with this packet and others in the capture, I could find out much more information about this share, such as the access controls used on the share, the authentication protocol used during the connection, and what was done during the time the connection was live. (We could look at the answers to the following: Were files created? Were files read? Were files written to?)

The point here, of course, is not to teach you packet analysis, but to make you aware of the amount of information available to anyone who can listen to communications occurring on your network.

Packet analyzers, monitors and sniffers can also be used for good. In addition to providing an excellent resource for troubleshooting network problems, the data can reveal attacks. If a malicious person is attempting an attack, capturing and analyzing packets on your network can help you determine what is occurring or what has occurred. An experienced person can also determine exactly what the attacker was attempting and whether or not she was successful. This is known as *intrusion detection*.

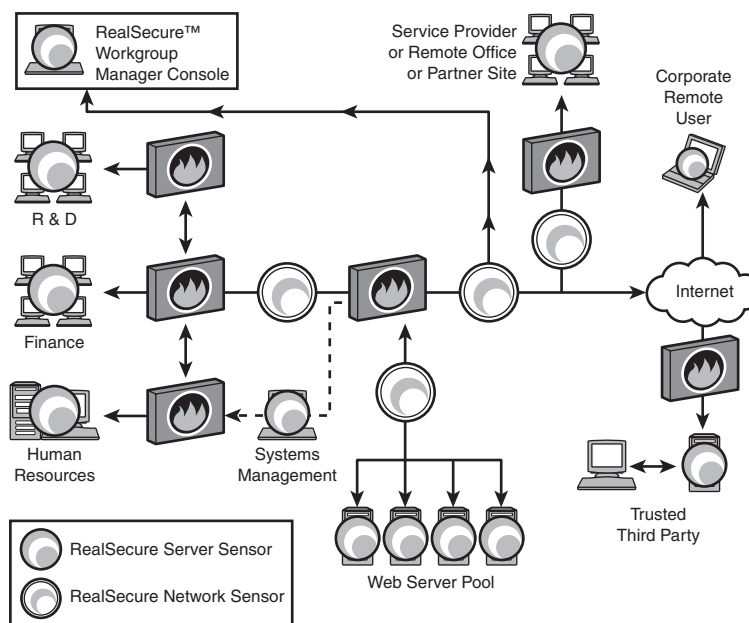
Intrusion detection is accomplished by extracting data and by the recognition of traffic and traffic patterns. Trained individuals can take the raw data produced by a network sniffer and deduce what is happening. Likewise, modern intrusion detection systems (IDSs) and applications attempt to take this knowledge and provide automatic alerting and even action based on programmatic analyses of events and the discovery of inappropriate, unusual, or incorrect activity. Some IDSs also use information from computer logs. Many IDS products exist. In fact, David Sobirey lists over 90 intrusion detection systems on his Web site at <http://www-rnks.informatik.tu-cottbus.de/~sobirey/ids.html>. Some of the more recognizable commercial products include BlackICE (http://www.iss.net/products_services/hsoffice_protection/), Cisco Secure IDS (http://www.cisco.com/warp/public/cc/pd/sqsw/sqidsz/prodlit/netra_ds.htm), eTrust Intrusion Detection (<http://www3.ca.com/Solutions/Product.asp?ID=163>), Network Flight Recorder (NFR) (<http://www.nfr.net/>), Real Secure (<http://www.iss.net/>), Shadow (<http://www.nswc.navy.mil/ISSEC/CID/>), and more than one free product, such as Snort available at www.snort.org.

Two types of IDSs exist: host and network. A *host-based IDS* requires loading software on the host machine. The software listens to traffic coming to and going from its host machine. It can also take advantage of information in the computer's logs and monitor the integrity of the file system for a broader picture of changes and attempted changes that might mean an intrusion attempt is in process or has occurred. To be effective, the host IDS software should be loaded on every computer. Host intrusion detection systems are considered more effective in detecting insider-based attacks.

A *network-based IDS* analyzes all traffic on the network. A central management station usually manages the information gathered by the host and network IDSs. Figure 7.3 diagrams this concept. In the figure, you can see the multiple RealSecure server sensors and the Manager Console.

FIGURE 7.3

Host sensors are located on all servers. This illustration is the copyrighted material of Internet Security Systems, Inc., reprinted by permission.



An example of a host-based intrusion detection system is host wrapper packages such as TCPWrappers for Unix, which are available from multiple sources on the Internet as a free download, including <http://coast.cs.purdue.edu/pub/tools/unix> and http://www.phys.ufl.edu/docs/system/public_domain/tcpwrapper.html. Nuke Nabber for Windows can be found at <http://www.dynaso1.com/puppet/nukenabber.html>. Other host-based systems are available as part of a personal firewall, such as BlackICE (http://www.iss.net/solutions/home_office/) or WRQ's AtGuard (<http://www.atguard.com>). For the large distributed network, companies such as Internet Security Systems offer host-based intrusion detection agents (for example, Real Secure Server Sensor, Real Secure's Desktop Protection).

NOTE

Intrusion Detection Resources For a fascinating, technical excursion into the knowledge provided by manual inspection of packet captures, use Stephen Northcutt's book, *Network Intrusion Detection, an Analyst's Handbook*. Windows 2000-specific traces are provided in sections of Thomas Lee's book, *Microsoft Windows 2000 TCP/IP Protocols and Services Technical Reference*.

In contrast, network-based systems gather information entirely by listening on the network. Different solutions are available including appliance style, such as RealSecure for Nokia and Cisco Secure IDS (http://www.cisco.com/warp/public/cc/pd/sqsw/sqidsz/prodlit/netra_ds.htm). Solutions are also available as software to be loaded on a computer attached to the network for this purpose, such as RealSecure, Computer Associates eTrust Intrusion Detection (<http://www3.ca.com/Solutions/Product.asp?ID=163>) and many more.

Both network- and host-based intrusion detection systems are based on *attack signature recognition* (the matching of known attack patterns with incoming data) and must be tuned and updated. A good IDS will provide an update service so that new attack signatures can be added. One of the tuning mechanisms is the capability to set the number of errors or instances of unusual activity that will cause an alarm. This is called *setting the clipping level*. For example, any system exposed to the Internet is subject to random port scans; in fact, it has become so common that we can almost liken it to the existence of background radiation. If an IDS were to alarm on every scan, not much other work would get done, and administrators would probably begin to ignore alarms. Setting a clipping level at some number of scans over the ordinary will provide warning when the normal background scanning has risen—perhaps indicating an attack directed at the network.

In addition to packet inspection-based intrusion detection systems, many products such as AXENT, Tripwire, and Cybersafe, provide unique host-based functions such as monitoring file system integrity and recognition of user access changes.

Penetration Testing Techniques

To defend the network against attack, you should not only be aware of generalized system hardening techniques, but you should also understand typical penetration (pen) testing techniques. In other words, you must study common scenarios to obtain information about the network and common attack techniques. Please note: I did not say you should use this knowledge to attack a system. Although pen testing of your own network can gain valuable insight into its vulnerabilities so that you can patch them, even this type of attack should not be done without written permission from the highest level of management possible, and it must be carried out by experienced personnel. Penetration testing is by definition intrusive, and some tools can harm systems. The object of *ethical hacking*, or the use of hacking tools to find vulnerabilities and patch them, is to secure networks, not to destroy them.

The following scenario discussion is purely fictional but representative of a common attack plan. An overview of the steps to take for penetration testing is listed in Step By Step 7.1.

NOTE

Clipping Level Is Not Just for IDSs

Clipping level is a useful concept for more than IDSs and often indicates the level at which errors become more than accidental. Many common occurrences can be an indication of an attack, or the result of human error. When is a failed logon an indication of password guessing or cracking in progress? How many visits to porn sites should be considered a violation of acceptable use rules? Can a single attempt to read a sensitive file be a reason for investigation? All these occurrences can be the results of simple errors.

If employees John and Sally have trouble remembering their passwords, but most other employees do not have this problem, and dozens of failed logons are recorded in the logs, something is obviously going on. A number of pornographic sites delight in having domain names similar to popular sites. Occasional hits on these sites should not result in an arrest warrant. However, if a particular file contains sensitive information and is difficult to access, even a single failed read can be cause for alarm. All these items and more can have dual meanings. Setting a clipping level will help you avoid over-reaction.

Setting a clipping level has a downside as well. Intruders know about clipping levels and will seek to slowly attack your system, hoping to remain beneath your “radar” to eventually break into your system.

STEP BY STEP

7.1 Penetration Testing

1. Determine the target. If your purpose is to gain notoriety, pick a very large and very public organization with a name that everyone recognizes. However, because that company probably has the best intrusion detection and security defenses and most knowledgeable administrators, you might want to pick a large company, not necessarily the biggest or most well known.
2. Footprint or profile. If possible, plant someone inside the targeted company and use social engineering techniques to obtain insider information. Low-level employees, such as janitors, guards, and other service personnel can plant bugs, steal documents, and social engineer information. Perhaps they can shoulder-surf and memorize a password being typed in or find passwords pasted to monitors or under keyboards.

In addition, use the Internet and other publicly available information (such as newspapers and magazine articles) about the company and its computing systems. Often the company will publish an amazing amount of useful information on its Web site, such as the location of data centers, new processing systems in place, and the names of software programs used. The Web site can indicate which type of Web server is in use. Searching the company's product information can also reveal information. If the company develops software products for IBM's AIX, it's a pretty sure bet that a large percentage of their internal servers, and maybe their Web server, are AIX as well. White papers, success stories, and partner lists can reveal what products your target is using and even the cities where they are deployed. SEC databases, employee profiles, and Usenet membership can provide useful information.

3. Enumerate the network. Common tools, many of them with legitimate uses, are readily available to enumerate (or map out machine names, IP addresses, and services) the network. Traceroute (or tracert in Windows) traces the path taken across the Internet to the Web server. This information can provide the name of the Internet service provider (ISP).

Then, locate the domain names registered to the company by using search engines on the Internet, reading articles about the company, and so forth. Run the whois tool on the Network Solutions (now owned by Verisign) www.netso1.com to learn information about the location of the site and possibly the name of the Web site administrator and the IP address of the DNS servers that have information on this site. Figure 7.4 shows the result of a whois inquiry.

Knowledge of DNS servers is also important as it may lead to the IP addresses of other computers that are part of this network. If the DNS servers are not properly protected, you may be able to obtain a full listing of Internet-facing (directly connected to the Internet) computers that belong to the company.

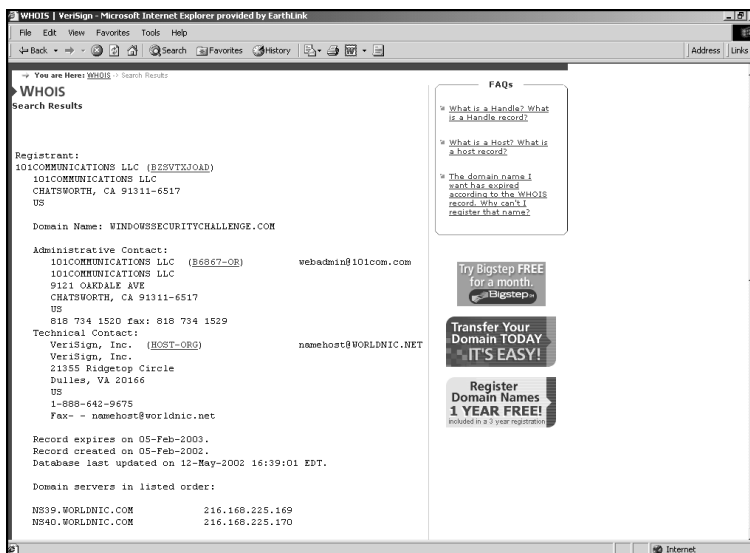


FIGURE 7.4

Using Whois to find the IP address of the Web server.

continues

FIGURE 7.5

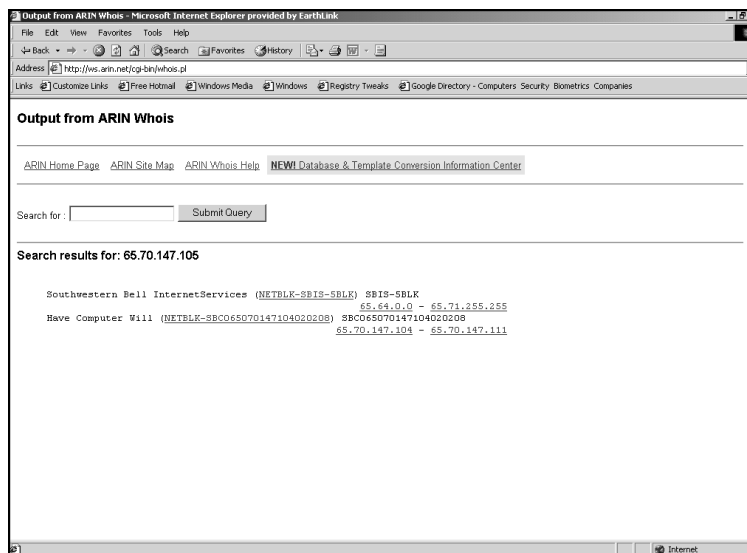
Using ARIN Whois to enumerate the network.

NOTE

UDP Scanning Although most port scanning looks for TCP ports, UDP port scanning can be used with varying success. UDP scanning can result in a false positive when access control is blocking UDP.

NOTE

Enumeration Elaboration In addition to knowledge of publicly accessible computers, information on internal computers can be found using a *war dialer*, a special program that dials a range of phone numbers and reports back on those where a modem answers. Modems represent remote access services. They might be sophisticated banks of modems tied to equally sophisticated services armed with strong authentication and authorization, or they might equally provide access to a single workstation that just happens to belong to the network administrator. Either provides an additional attack vector. Some war dialers are quite sophisticated and will also test connections, run password crackers, or attempt to determine services running on the computers they reach. War dialers, such as blue beep, phone tag, and PBX Scanner, can be found at <http://packetstorm.decepticons.org/wardialers/>. Commercial war dialing programs, such as Sandtrap (<http://www.sandstorm.net>), that can dial multiple phone lines simultaneously are developed to assist in vulnerability testing and are also available.

*continues**continued*

4. Scan and enumerate services. Armed with the IP addresses of publicly available servers, you can now use various scanning tools to learn more. For example, Ping sweeps return echoes from live hosts—you know a computer is using that address. Some systems, such as firewalls, can be configured to block these replies, so your ping sweep will not find them even though they might be live and operational on the Internet. Port scanning tools attempt to determine the services running on the computers. Although you can use crafted packets, common tools available, such as NMAP by Fyodor at <http://www.insecure.org/nmap/>, can do ping sweeps, port scans, and OS detection. A Windows version of the product is available from <http://www.eeye.com/html/Research/Tools/nmapnt.html>. Other tools include NetScan Tools Pro 2001, Superscan, and Fscan from Foundstone (www.foundstone.com). Knowledge of the services running on the computer is useful because attacks are designed to take advantage of a known vulnerability in a particular service. The Code Red, for example, attacked port 80, the port used for Web services. If our scans locate

open ports, we can choose our attack tools.

5. OS Enumeration. Scanning can also find other information. Many services when they receive a connection request, issue a banner, or string of information. It can include the name of the service as well as the operating system version. Port scanners can return this information. In addition, a telnet client can be used. Directing the telnet client at a port commonly used for a particular service will usually display the returned banner.

Another tool that can determine the OS (and other information) is netcat (Unix and Windows versions are available from <http://www.atstake.com/research/tools/#forensic>). Knowing the name of the Web server or operating system allows a directed attack using knowledge of a vulnerability associated with that operating system or service. The process of seeking this information from the information provided is banner grabbing. OS fingerprinting can also be accomplished due to subtle differences in the responses of different TCP/IP implementation, and by common Web page extensions. The existence of a page called Mywebpag.asp, for example would identify a Web page on a Microsoft Internet Information Server because the .asp extension follows the page name.

6. Penetration test. The final phase involves an attack against a particular machine. The tool used or code written depends on knowledge gained to this point. The goal is to obtain privileges (the rights to do something) and access rights on the machine. The use of a tool or custom script to increase the user's rights to that of administrator or root is called an *elevated privileges attack*. After you have gained some level of control on a single system, use this to further footprint the rest of the victim network. What can be learned about the other systems? Does this system have special connectivity to those systems? What other networks is it connected to? If this system is the Web server, is there a credit card database or other interesting data stored locally? Are there links to databases?

NOTE

continued

Use the American Registry for Internet Numbers (ARIN) (<http://www.arin.net/whois/index.html>) whois tool to determine the IP address block assignment for a company (see Figure 7.5). Entering one known IP address in the whois tool returns the range of addresses assigned to a particular domain. These will be routable Internet addresses, which means if they are assigned to an Internet-connected computer, they are reachable and attackable from the Internet. Note that an address might be assigned, but might not represent an Internet-connected computer. You might also have to use Arin's counterpart to obtain this kind of information.

For Europe, use

<http://www.ripe.net/ripenncc/pub-services/db/whois/whois.html>.

For Asia, use

<http://www.apnic.net/search/index.html>.

Samspade, a free tool, can be downloaded from <http://samspade.org> and used to footprint the network.

NOTE

How to Hack a Bank It's often helpful to imagine how to put the attack steps into play. A useful description of an imagined hack in which millions or billions of dollars are stolen from a bank can be found at www.infowar.com/hacker/00/hack_052200a_j.shtml. The document outlines how a team of well-heeled and knowledgeable hackers could succeed in this attack. The step-by-step account is logical, believable, and chilling.

DEVELOPING COUNTERMEASURES TO THREATS

Define threats and countermeasures.

The way to eliminate or mitigate risk is to develop and follow countermeasures for each identified threat to information systems. It sounds so simple, doesn't it? What complicates this seemingly straightforward approach to security is the existence of multiple threats and their continually changing nature. Threats that yesterday were considered unlikely are now possible. Some threats seem to have little risk, and therefore companies are less likely to apply the countermeasure if costly or inconvenient. Not all that long ago, although airlines recognized the threat of airplane hijacking, they felt the inconvenience of applying extra countermeasures outweighed the slight risk. 9/11 changed that, and since then we have seen increased vigilance and security measures at all U.S. airports.

Risk analysis determines which threats require development and implementation of countermeasures.

Risk Analysis

The process of risk analysis is used to determine whether threats to systems will result in damage. An analysis of vulnerability and possibility determines how great the risk might be. Risk analysis often results in a ranking of threats from those most likely to those least likely to cause damage. This ranking then determines the expenditure of resources including money and staff in a direct proportion to the level of risk. Two methods are used:

- ◆ **Quantitative risk analysis**—Involves multiplying the probability that an event will occur times the monetary loss. Typical formulas used are Annual Loss Expectancy (ALE) and Expected Annual Cost (EAC). This process is difficult (because it's difficult to figure out what the reliable probabilities are) and time-consuming. Automated commercial products, which do the calculations for you and even recommend risks to quantify, are available.

- ◆ **Qualitative risk analysis**—Uses the estimated loss and evaluates each threat by looking at specific system vulnerabilities and noting the countermeasures (controls) already applied. It is often referred to as the “gut-feel” analysis protocol because the result is often tempered by the collective experience of the participants, not by statistics.

Threats

Risk analysis is conducted on and countermeasures are developed for perceived threats. Table 7.3 lists common information system threats and describes examples.

TABLE 7.3

COMMON INFORMATION SYSTEM THREATS

<i>Threat</i>	<i>Notes</i>	<i>Example</i>
Errors	Incorrect passwords configuration.	Default, well-known are not changed.
Omission	Patches are not applied.	Patches for IIS were not applied and many IIS servers were infected with Code Red.
Fraud	Company assets are obtained by misrepresentation, or modification of information.	Paycheck amounts increased by claiming overtime hours not worked, customer records stolen, or software taken by employees for home use.
Misuse of information	Sensitive, private information is used for personal gain.	Use of earnings knowledge used to buy or sell shares (insider trading).
Employee sabotage	Employee uses knowledge of company operations and systems to destroy or damage.	Time bombed code loaded on servers by administrator destroys data the day after the employee is fired.
Ignoring policy	Employees know the rules but do not obey.	Accidents caused by not following safety rules. Accidental destruction of data backup by leaving tapes in the trunk of a parked car during a summer heat wave when policy states immediate transport in air conditioned vehicle.
Physical accidents	These are the result of physical circumstances as opposed to system malfunction, or inadvertent misuse of the system.	Electric shock, moving parts of printers.
Software malfunction	Bugs or security vulnerabilities.	Buffer overflow causes reboot or leaves the system open to compromise.

continues

410 Part I EXAM PREPARATION

TABLE 7.3 *continued***COMMON INFORMATION SYSTEM THREATS**

<i>Threat</i>	<i>Notes</i>	<i>Example</i>
Loss of resources	Destruction of data center in full or part.	Fire, flood, storm, bomb, or explosion.
Loss of infrastructure	Malfunction of equipment.	A router or switch dies.
Hackers and crackers	Attack on systems.	Loss of data, loss of reputation, and destruction of systems.
Espionage	Spies from another company join yours, or pay your employees to provide internal information.	Soft-drink formula is stolen from database by employee and sold to competition.
Malicious code	Code is run on system with undesirable results.	Code Red, Nimda, I Love You, and so forth.

In the mainframe world, several operations personnel were required. The job of each person was carefully defined, and extensive work has been done on the threat model for the operations group. This is not the case for modern PC-based distributed systems. An example of mainframe/operations threats and countermeasures is provided in Table 7.4.

TABLE 7.4**EMPLOYEE JOB DUTIES, ACCESS LEVEL, AND RISK**

<i>Job</i>	<i>Description</i>	<i>Access Level</i>	<i>Risk</i>
Computer operator	Do backups, run jobs, mount tapes, load paper in printers, record, report problems, operate with devices, software products, system performance metering, heat control, humidity controls	Console, tape/disk drives, printers, operations documentation, problem/change management system	Gains access to production data files, production maintenance, and job control, program documentation; turns off logging (can lose audit trail) potential loss of system records due to not enough roomon media
Operations analyst	Analyzes computer memory and hardware requirements; estimates use of disk and tape, performance; advises on operations documentation; establishes backup, recovery procedures; monitors service level agreements; installs new hardware and telecommunications; replaces obsolete items, and troubleshoots	Test files, operation documentation, system performance reports	Access to production data files and production application programs
Job control analyst	Job control language; assists application programmers; reviews production problems using problem change management process; tests and implements new features; and assists in product troubleshooting	Test job control files, job scheduling files, operations documentation, problem/change management system	Access to production data files, application programs, and job control files

<i>Job</i>	<i>Description</i>	<i>Access Level</i>	<i>Risk</i>
Production scheduler	Plans, creates, and coordinates computer processing schedules for production jobs and job streams; consults with end users and application programmers concerning production schedules; completes ad hoc jobs; reviews results in comparison to planned schedules; and updates and issues monthly billing schedules	Job scheduling files, operations documentation, problem/change management system	Access to production files, data files, production application programs, and job control files
Production control analyst	Printing, balancing, distribution of reports and records, manages printer, burster, and decollator, balances required reports, assists production scheduler, and performs inventory counts and computer supplies	Computer equipment, supplies and reports, and problem/change management system	Delivers reports to wrong individuals, theft of supplies
Tape librarian	Collects input tapes; sends/receives tapes from off-site storage; maintains tapes and cartridges; ensures adequate supply, tape storage, and vault; ensures critical backup; pulls historical files and stores at local tape vault or ships to offsite location; maintains logs; and controls physical inventory of tape library	Automated tape library, problem change management system	Production data files, application programs, and job control files

Countermeasures

Countermeasures can include general system and network hardening steps, or they can represent special efforts directed at specific threats or specific vulnerabilities. Many hardening steps are not system specific, such as disabling unnecessary services, patching systems, maintaining a strong password policy, requiring strong authentication, logging and analyzing logs, and training users in security awareness. Others are specific to the operating system, Web server, service, application or device.

Checklists for hardening systems and applications can be found on the Web sites of the company that produces the software or hardware. Many vendor-neutral locations for checklists on multiple operating systems exist, such as SANS, located at www.sans.org. The National Security Agency (www.nsa.gov) also provides numerous computer security checklists on Windows, Cisco routers, and countermeasures to malicious email content.

NOTE

A New Threat Model New situations require new terminology. Today's wide adoption of the Internet and its widespread knowledgebase free for the asking; the existence of malicious code which duplicates and passes itself on to its next victim; and the existence of prewritten attack scripts that every clueless person can run all mean more attacks on companies and individuals of every kind. John Kindervag, whose opinion is explained at www.osopinion.com/perl/printer/17692, has developed a new taxonomy of threats for the modern world. They are

continues

NOTE

continued

- **Strategic attack**—In this attack, you or your company is picked as the target.
- **Collateral attack**—This attack is directed at some other company or individual but gets you and/or yours as well.
- **Nuke attack**—You or your company suffers simply because you're connected to the Internet. Worms and viruses form the majority of these attacks and can affect everyone on the Internet.
- **Random attack**—In these attacks, automated tools scan huge numbers of IP addresses looking for vulnerabilities. Of the multiple easier victims, you're selected.
- **Jump-point attack**—In these attacks, your computer is compromised and used to attack others.

Risk analysis should consider these threat models when calculating probability and possibility, and you should develop countermeasures for them as well. Two of these models, the collateral attack and the nuke attack, should be judged as probable for all companies and individuals with an Internet presence. Those attacks that focus on a specific target might be less of a threat to those with a low profile. Attackers will go after the big guys. However, this does not mean that countermeasures should not be developed; it just means that the small business might not need to spend the money that the larger business cannot afford not to invest. Countermeasures for these attacks are complex and consist of the standard security practices as well as special and extreme measures to protect the more valu-

Establishing Countermeasures for Employee-Related Threats

Many threats can be classified as employee related, and developing strong countermeasures can therefore reduce the risk associated with these threats. Following are some of the mitigating procedures you can apply to reduce the risk of threats:

- ◆ **Provide clear definition of authority**—Define the responsibilities of all positions and identify who reports to whom and who supervises whom. By indicating who should be doing what, you allow all employees to question someone who seems to be doing something they should not. Employees cannot determine what is questionable activity if they do not know what is proper. Knowing who has authority is paramount.
- ◆ **Structure along functional lines**—Employees can be prevented from entering work areas they don't belong in. Many system compromises, for example, are the result of physical access to servers. If data center employees know who should or should not be in the data center, they can comfortably challenge strangers or report them to security and perhaps prevent an attack.
- ◆ **Ensure that any type of fraudulent behavior requires collaboration of two or more individuals**—This discourages fraud and makes it harder. It also makes it more detectable, as surprise audits or unexpected events can reveal suspicious activity. If one person, acting alone, can subvert the system, it can be hard to detect because only he knows what he did. However, if he must obtain cooperation from others, multiple people are involved. A misstep by one of them can expose the activity. Separation of duties and structuring along functional lines help enforce this requirement.
- ◆ **Separate job functions when combining them provides too much control**—In small shops it is often difficult to have one employee for each job function; therefore, many employees do more than one job. Careful attention to job combinations reduces opportunities for misuse of power, accidental poor practices, and fraud. Separate systems analysis duties from programming, systems development from systems maintenance, and operations from development.

- ◆ **Rotate people within their own areas**—Software maintenance employees, for example, can be moved from the responsibility for one application to another. Operations people can be required to rotate shifts. These activities promote cross training, ensuring the ability to recover from disaster or maintain function during shortages. Rotation also prevents collusion because no individual remains in charge of an area or has authority at the same time.
- ◆ **Prevent family members from holding jobs in areas which you would not combine into one person's responsibilities**—If, for example, an employee works in operations, his wife should not be hired to work in that area, or be assigned to the programming staff.
- ◆ **Provide clean, accurate, detailed job descriptions**—If an employee knows what she is supposed to do, she is less likely to make an honest mistake that results in system compromise or fraud. For example, if Mary orders hundreds of thousands of dollars in computing equipment and has them shipped to a noncompany address, she cannot claim that she did not know she was not authorized to do so—if her job description specifies a purchasing range below this and specific instructions to only ship to addresses in the company data center address database. Better still, programmers working on the purchasing application can code in purchasing rules that will prevent Mary from overstepping her responsibilities.
- ◆ **Include as part of every employee performance review, evaluation, and consideration for raise and promotion the employee's observance of security practices**—An interesting implementation of this might be occurring in software development companies. In early 2002 Microsoft announced that employee bonuses would be pegged to the production of secure code. Although details are sparse, and it is too early to determine the results of such a policy, it certainly could be a strong countermeasure if enforced.
- ◆ **Provide annual training for all employees**—Training should include review of security objectives and policies, practices of the organization, and the employee's responsibilities during disaster. Employees in sensitive positions should receive more comprehensive and more frequent training.

- ◆ **Encourage IT security to work with other security specialists, such as plant and physical security**—When specialists work together they can strengthen overall security.
- ◆ **Maintain a standards manual and enforce the standards.**
- ◆ **Require vacations be taken and require that they be taken contiguously**—When regular employees are on vacation, others must do their jobs. This often allows discovery of fraudulent or suspicious activity. It also can uncover simple errors of omission or practices that produce vulnerabilities. If the administrator has not been updating patches or reviewing the logs, his substitute can quickly discover this while the administrator is on vacation.
- ◆ **Require sophisticated access controls at the entrances to sensitive areas and systems**—Guards, ID cards, smart cards, and biometrics can prevent improper access to physical areas. Smart cards and biometrics can also control access to data systems.

Including Countermeasures in Hiring and Firing/Exit Practices

Countermeasures should also be a part of hiring and exit practices. Background checks of individuals with responsibilities for data and data systems can avert many problems. Fraud investigations often turn up evidence of fraud, theft, sabotage, and misuse of information by employees at previous jobs. After an employee is hired, monitoring behavior during probational periods can also assist in uncovering questionable behavior.

Information that should be required and actions that can be taken when candidates apply for jobs that give access to sensitive data or responsibility for administration of data systems include

- ◆ **Requiring business and personal references**—Many companies prefer personal references from professionals such as lawyers, doctors, dentists, and clergy.

- ◆ **Making employment contingent upon receiving a reference from the candidate's current employer**—Understandably, this is not always available during the interview process, but it should be requested from the former employer before the employee starts work.
- ◆ **Checking public records, including court records, marital record, educational record, military record, law enforcement records, public documents and credit bureaus.**
- ◆ **Requiring drug testing.**
- ◆ **Considering insurance and bonding**—A surety bond reimburses a company for loss due to theft of specific assets and fraud.
- ◆ **Looking for conflicts of interest**—Has the candidate received fees from vendors for obtaining business?

Investigation of part-time employees might be necessary as well, depending on the nature of the job and the length of employment.

When employees leave the company, whether they resign or are fired, strong countermeasures should be applied. An exit procedure should be defined that includes a checklist of duties. Items of importance are the collection of keys, ID cards, and other company materials and the changing of locks, passwords, and other access codes.

Gruntling Program

It's commonly said that disgruntled employees are responsible for much employee fraud, destruction of data, and other malfeasance. More than one commentator has said in reply, "You need a gruntling program then."

Often, employees who sabotage are quoted as saying that no one cared, that the company treated them like things, not people. It's clear that a policy that promotes employee satisfaction and removes the common causes of disgruntlement is long overdue. Consider it a countermeasure to employee-related threats. Here are some ideas that might work:

NOTE

Other Than IT Employees? Other employees can directly impact the security of information systems. Vendor employees, air conditioning, maintenance engineers, and building personnel all have contact with equipment and are provided entry into protected areas. Good countermeasures are to require bonding, not to allow free access (instead require service orders), and to require identification and signing in and out. Observe to ensure personnel only access required equipment and only enter required areas of restricted areas. Always escort people into secure areas, and never leave them unattended.

NOTE

Double Take! Is the IT department at your company notified of employee exits? Frankly, this is a huge problem. IT departments should be notified, but frequently they are not. Unless an IT employee validates user accounts on a regular basis, an audit might uncover numerous accounts still enabled years after employees have left the company. At one account where I recently assisted in an audit, we found over 1,000 accounts that had not been used in over six months!

Countermeasures include setting expiration times for accounts and scanning logon records to find accounts that have not been used in several months. Automated utilities exist to assist in finding this information.

NOTE

Comfortable Seats Move Products

Faster Long-distance truckers must sit for long periods of time in less-than-armchair comfort. Many years ago, truckers at a major Midwest trucking company demanded expensive, comfortable seats for company trucks. Is this a gripe or a solvable grievance? After much research, testing, and study the truckers got their seats. Why? Because the tests proved that when trucks were fitted with the more comfortable seats, truckers drove for longer periods of time with fewer accidents. Trucks reached their destinations quicker. Analysis determined that the more expensive seats were actually cheaper in the long run. If the truckers' griping had not been heard, the study would not have been conducted, and the company would not have found that the resolution to the complaint was also good for the business.

- ◆ **Respect employees and consider individual situations.**
- ◆ **Consider morale-building programs**—Develop pride in company products, philosophy, and attitude. When morale is poor and employees don't work as a team they might not consider preventing other employees from breaking the rules, or they might ignore employees who do break the rules.
- ◆ **Provide security training on an annual basis**—Refreshing memory and deepening understanding for security practices can go a long way toward obtaining employee buy-in. Employees are less likely to grumble about security practices that seem to restrict them in their jobs if they understand these practices help them keep their jobs.
- ◆ **Provide professional development opportunities**—Encourage and pay for job-related skills training and for training that will help the employee advance in her career path.
- ◆ **Provide rewards for good behavior, such as bonuses and other recognition of accomplishments**—Pay special attention to rewards for struggling with frequent and unrealistic deadlines, one of the prime stress factors in IT departments.
- ◆ **Increase communications through staff meetings, group meetings, and discussions in which employees can air gripes and grievances**—Then do something about them (the gripes, not the people airing them). Gripes might be the result of misunderstanding, or they might be about things that cannot be changed. The important thing is that open discussion can mitigate the hostility that can result from ignoring the problem.

Countermeasures for Common Internet-Based Threats

In Step By Step 7.1 an over-the-Internet attack procedure was outlined. To mitigate the risk of this type of threat, the following countermeasures can be applied:

- ◆ **Footprinting/enumerating the network**—Most information gained here is public knowledge. You can, however, obscure some information. For example, ensure that contact information listed for domain registration is general; in other words it does not contain a real individual's name.
- ◆ **Scanning/enumerating services**—Block all unnecessary inbound and outbound ports. This can be done at routers and should be done at firewalls. Inbound ports are blocked to prevent attacks that take advantage of vulnerabilities in the related services. Outbound ports are blocked so that an attack originating from within the network will not be passed outside the network. Even when ports are blocked, unnecessary services can provide vectors for attack. The wily hacker will use a port redirection tool to attack a known vulnerability in a service by using a port that is not blocked. If services are not used, they should be disabled. In some cases it is possible to use IPSec to filter or block access to all ports except those required. *IPSec* is a security protocol that can be built into or added to the TCP/IP networking software on a computer.
- ◆ **OS enumeration**—Because many hints are found in banners, or notices returned when enquires are made, where possible change or eliminate the banner presented by services.
- ◆ **Penetration test**—Become knowledgeable of the tools and tests that hackers use. Develop or find tools that are countermeasures to these tools and methods. Continually research vulnerabilities and develop countermeasures. Apply patches and configure systems appropriately. Use intrusion detection/prevention methods and programs.

Countermeasures to Physical Threats

Physical threats also have related countermeasures that can mitigate or eliminate their risk. These include the following:

- ◆ Don't build near explosion hazards; also, don't locate a data center near any explosives. In addition, diesel-powered generators should not be located near the data center.

- ◆ To avoid windstorm damage, don't have exterior windows, and provide protection from possible falling trees or manmade structures such as towers.
- ◆ Don't place the data center on lower floors. Break-ins occur more often on lower floors.
- ◆ Do not externally label data center locations or advertise in phone books, Web sites, and so forth.
- ◆ Avoid basement locations. Water damage can result from flooding. Use watertight seals and reroute pipes and conduits away from the data center if possible.
- ◆ Don't place media storage areas/vaults near flammable or explosive material, and don't place them near compressors, water, and gas tanks.
- ◆ Subdivide rooms with firewalls or man traps, and keep fire doors closed.
- ◆ Use noncombustible building materials.
- ◆ Store paper media separate from equipment.

THE ROLE OF ADMINISTRATIVE MANAGEMENT

Define the role of Administrative management in operations security.

Administrative management, the management of all things administrative, can serve a critical role in operations security. Managers must concern themselves with legal compliance, risk management, and fiduciary (monetary) responsibility. These are impacted by operations security. In addition, management plays a key role in promoting education on security, overseeing compliance, participating in policy-making and enforcement, ensuring cross-departmental involvement and approving funding.

In fact, administrative management's role is tied so closely to operations security that their lack of attention to security represents a threat.

What if management flaunts controls? If plant manager Bob insists on having root authority on the production server, could he not manage to defraud the company by changing production numbers and selling product on the side? If office manager Mary can change the configuration on her desktop, won't she be tempted to set up her office PC so she can access it from any browser on the Internet? If sales manager Peter has permission throughout the departmental sever, could he not accidentally erase data files? Sure they could. What's more, they could be more easily socially engineered to relinquish passwords for their privileged accounts than someone properly trained in the consequences of doing so.

What if management does not fund security? This is not only a very real threat; for many companies it is a reality. Funding for security products, training, and practice is often shortsighted. In the past, management has often underfunded security by hiding behind the shield of probability. Admittedly, in the past the probability of an attack on most business networks was highly improbable based on the effort required to do so, the fact that few businesses were well integrated with the Internet, and, perhaps, even a common belief that doing so was wrong. These factors have changed. Today, businesses need increased funding for security measures and a continued commitment on the part of management.

Management must also take a role in information security. They should be involved in the definition of its scope and in the preparation of a statement of the results to be achieved. Security objectives should be a part of general organizational objectives. Management can help coordinate security activities both in IT departments and in other areas of the company. Management can monitor the process, obtain feedback, monitor results, obtain resources, promote inter-departmental programs, develop relationships, obtain money and facilities, and assign responsibilities and authority to individuals.

Security managers have a more direct role. They are often part of an IT department but might belong to a separate security department. This is a growing trend and makes sense, as it fulfills the separation of duties principle. Typical security manager job titles include Information Security Administrator, Computer Security Manager, Security Information System Officer, and Security Officer.

Job requirements include managerial and technical talents. A security manager should be able to evaluate technology solutions; promote security awareness; initiate technical, managerial, and people solutions to problems; and sell security concepts to every strata of the organization. In addition to general security knowledge she should have deep knowledge of auditing, internal control, risk analysis and industry-specific security issues. Security managers are often required to seek certification in technology and security management. A number of vendor specific security certifications are reviewed at <http://certcities.com/certs/other/>. Table 7.5 lists specific security certifications for managers.

TABLE 7.5**CERTIFICATIONS FOR SECURITY MANAGERS**

<i>Title</i>	<i>Initials</i>	<i>Manages Certification</i>
Certified Information System Security Professional	CISSP	(ISC) ² www.isc2.org
Certified Information Systems Auditor	CISA	Information Systems Audit and Control Association www.isaca.org
Various vendor-neutral certifications on security management and technical areas	Various	SANS www.sans.org

CONCEPTS AND BEST PRACTICES

Define operations security concepts and describe operations security best practices.

- **Explain antivirus controls and provisions for secure email.**
- **Explain the purpose of data backup.**
- **Detail how sensitive information and media should be handled.**
- **Describe how media should be handled.**

Throughout this chapter many security principles have been discussed. Information has been provided on how the practices of least privilege, separation of duties, and change management can improve security and reduce the risk of fraud and accidental loss of data or data integrity. However, many other operations best practices contribute to the stability and security of information. Some of them are discussed in other domains. Legal issues, for example, such as legal requirements; the standards of due care/due diligence; and record retention, privacy, and protection are discussed in the legal domain. Data backup is discussed in the Disaster Recovery and Business Continuity Domain. Additional operations security concepts and best practices are

- ◆ Privileged operation functions
- ◆ Email security including antivirus controls
- ◆ Protecting sensitive information and media
- ◆ Change management

Each of these is discussed in the following sections.

Privileged Operation Functions

Privileged operations are system commands and parameters and the configuration commands and activities for any device that handles information or controls the transmission of data on the network. This includes tape systems, external drives, communication devices, and infrastructure (router, switches, and so forth), as well as computers.

In the past, misuse of these commands was prevented by tight control over the knowledge of these commands and their parameters, as well as by restricting their use and protecting the consoles and devices necessary to issue them. The practice of control by information obscurity is no longer followed; indeed it would be impossible to do so. Two factors are responsible.

First, the post-mainframe, post-Unix world has a tendency to empower the user at the expense of protecting the operating system. This, of course, is the result of the need for individual management of the multiple systems that were brought into the organization by the users themselves. At first, no support was provided—what can you expect?

Second, the number of computers, the jobs that they do, and of course, the infrastructure that supports them have exploded. The secretive, wizards-of-the-temple-of-IT method of knowledge transfer just does not scale. With the explosion of computers and their infiltration into every function of a modern society has come a proliferation of knowledge. Information is available from numerous sources including books, Web sites, colleges, and technical training programs. Unfortunately, the widespread availability of information means that even though it's easy to find someone who knows the *how to* of systems administration, it is difficult to find someone who also knows the *when to* and the *why*. We now have many systems administrators who know little about security or the impact of what they do.

The solution, like the problem has multiple parts. First, we must ensure that system commands and utilities are reserved for administrative use. Second, we must provide training and guidance for all administrators, in the why and wherefore of what they are charged to do. Finally, we must ensure that job interviews also stress this aspect and not just rely on technical competency.

IN THE FIELD

TRAINING CONSULTANTS IN SECURITY

It has been my privilege to train not only information system auditors but system administrators, help desk personnel, and IT consultants in the technical why and how of operating system and server application security. But never was the how-to-versus-why conundrum brought home more than during a week-long intensive training session with sixteen senior IT consultants. This class was tightly focused around installing and configuring secure email and included technical training on virtual private networks (VPNs), firewalls, mail servers, and public key infrastructure (PKI). During the class consultants participated in either of two teams representing two companies. After instruction and a series of labs on each technology, each team was instructed to implement a secure email system. They were given a list of goals, and appropriate hardware and software to complete the task.

Quite inadvertently, I had placed the majority of the more technically competent people on one team. They advanced more quickly on the assignment. At one point, however, it became obvious that the VPN solution was not working as advertised. The specifications of the project required the VPN to use a specific authentication algorithm and data encryption. The more experienced team was first to

call the vendor for assistance, not because they lacked technical savvy or willingness to solve the issue, but because they began working with the product and exhausted the possibilities sooner. The end result was a relaxation in project specification. The VPN, it seems, could not perform the required authentication protocol and could only use a much weaker process.

The less experienced team would not accept this solution. Instead, they investigated the authentication protocols available from the native operating system and found that one met the project specifications. They returned to me as the provider of the specification asking approval to use this solution instead of the other vendor products.

This difference in approach also was present in another scenario. Students were asked to appropriately configure the mail server for administration. Both groups were told they needed to provide administrative accounts that had authority to manage, troubleshoot, and maintain the mail server. Three possible privilege assignments existed: user, mail admin, and service account admin. The more experienced group accomplished this with one step; they gave the local Administrators group service account admin privileges. The other group created a mail server administration group and in addition to local Administrators group membership only gave them mail admin privileges. The difference is that although both groups accomplished the stated goal, the first group gave more privileges to more people than necessary. The second group restricted the ability to administer the server to a select group, not the entire group of operating system admins. They also correctly assigned only the privileges necessary for administration.

In project review, both teams discussed the issues. The more experienced group focused on getting the job done. The less experienced group focused on getting the job done right.

Understanding Antiviral Controls

No one would question the need for antiviral controls. Antivirus products are one of the few security-related products that usually are approved for purchase. People seem to realize the need for this type of protection. Why then, do we continue to hear that viruses and worms account for so much damage, congestion, and disruption?

Clearly, to purchase and install antiviral remedies is not enough. Mitigating the threat of virus infection takes technology; savvy administration; informed, cooperative users; and technical controls to make them work.

Medical analogies work well in a discussion of computer viruses. We call them viruses; our computers get infected; and we inoculate mail servers, file servers, firewalls, and desktop systems against the risk. Here's another parallel: Until every computer and every Internet or network connected device not only runs antiviral software, but keeps it updated, and until every administrator and user understands and follows a strong antiviral protocol, we won't be rid of any of them. So, what is an antiviral protocol?

Medical practitioners don't generally do just one thing to combat a disease. They don't just prescribe drugs, or perform surgery; instead each disease and each health concern has a strong management protocol that prevents reinfection as well as treats symptoms. Best practices for antiviral management need that as well. Five areas must be addressed:

- ◆ **Antiviral products must be installed on servers and desktops**—Specialized mail server versions of major antiviral products exist and should be used. All desktop systems must also have software installed.
- ◆ **Automatic, regular updating of both engine and patterns is a must at the server and desktop levels.**
- ◆ **Server side products should be configured to use additional features**—Blocking of executable attachments to email is one example of a server side feature. As a major entry point for viruses into the system, email server-based antiviral products can assist in protecting other systems but must be properly configured and tested to ensure that they work.
- ◆ **Attention should be paid to new viral/worm vectors**—All infections will not come from email or desktop systems. Any computer or device running instant messaging, Internet Relay Chat (IRC), and Personal Digital Assistants (PDA), or other wireless devices all can become infected and pass the infection on. Some malicious software will attack multiple entry points including Web servers, messaging, email, and software transfer.

- ◆ **All users should be trained to not accept defaults, to be proactive, and to resist social engineering techniques.**

In sum, mitigating the threat of virus attacks requires much more than simple installation of antivirus products. A solid program will pull together multiple defensive actions.

Protecting Sensitive Information and Media

Sensitive information is any information for which distribution should be managed, rather than available to the public. Often this information's availability to other than its intended audience results in problems.

Sensitive information has varying degrees of sensitivity. Customer lists, for example, should be available to company sales people and to those managing accounts payable but should not be published where competitors could obtain them. Information that might adversely affect the market value of company stock should not generally be available to any employee. On the other hand, the location of corporate headquarters or current product descriptions is information that belongs in the public domain. Military information security standards also have their own system of data classification, but the principle is the same. Sensitive data needs to be managed differently.

How should sensitive information be managed? Sensitive information and the media it is available on should be more carefully managed. Information, like all things, has a life cycle. It is created (purchased, discovered, developed), handled, stored, and finally destroyed. Each phase requires specialized handling. The phases are

- ◆ **Creation**—All data, however it is obtained, should immediately be classified and labeled. The labeling should indicate when it was obtained, its source, and an indication of its sensitivity level. Data that is stored and used electronically should also be identified electronically.

NOTE

Removing Data from RAM and ROM?

Clearing sensitive data from disks can be accomplished in several ways such as deletion and overwriting or degaussing (de-magnetizing).

Removing data from Random Access Memory (RAM) is usually done by clearing or by removing power. Data in Read Only Memory (ROM) is permanently stored.

- ◆ **Handling**—All data within the data center must be properly handled to assure viability and confidentiality. Protect media by keeping it in original packaging away from direct exposure to heat, sunlight, and electrical shock or damage from dropping. When necessary to transport media, it should always be moved directly from the computer room in air-conditioned vehicles. Media should be stabilized in the computer room for 24 hours before using it. Ensure that labeling accompanies the media, and ensure that changes in media location are recorded. A manual log usually serves to identify storage location and details the who, what, why, when, and where of any movement.
- ◆ **Storage**—Provide environmental controls such as the ideal temperature and humidity level and freedom from dust and dirt. Printers should not be stored in media storage areas. Printers increase the level of dust, and laser printers increase the level of ozone in the air. Ozone can cause changes to media. Care should be taken to ensure that air-conditioning intakes are not located where they can bring in diesel fumes from loading docks. *Positive air pressure*, or when air blows toward the door, will help to maintain a better environment. Schedule air conditioning maintenance for cooler times of the year. Wax and cleaning agents should not be used on computer room or storage area floors. The solvents, dust, and wax particles as well as the debris from the buffer can damage media and equipment. Air conditioning should not be shut down at night or on weekends. This causes temperature and humidity changes in media which can be damaging.
- ◆ **Cleaning**—Wax and cleaning agents should not be used in computer room or storage area floors.
- ◆ **Destruction**—When it is no longer necessary to maintain data, the data should be destroyed. Common practices include clearing and purging. Clearing removes data from media but does not take the extra steps that would prevent recovery of data if the media can be subjected to laboratory attacks such as strenuous forensic techniques. Clearing does prevent recovery of data using a *keyboard attack*, a technique that uses common system utilities or software. Clearing is adequate when the media will be reused on the same computer in a physically secure place, or when the data, which was removed, is not sensitive.

Purging takes a further step by preventing recovery even if the media is subjected to laboratory tests. Methods used include multiple overwrite of data, encryption, media destruction, and degaussing. *Degaussing* magnetically erases the disk contents. Destruction is via a metal destruction facility such as a smelter, or via pulverization, abrasion, incineration, or acid wash.

Change Management Control

Change management control is often described as a best practice for management of custom software development and maintenance. Computer operations should also institute a change management control system for IT infrastructure. The first step in the process should be to develop detailed documentation on the following:

- ◆ Network configuration
- ◆ Computer configuration
- ◆ System parameters and settings
- ◆ Application configuration
- ◆ Device configuration
- ◆ Locations for all computers, devices, media storage, and other parts of the infrastructure
- ◆ Job titles and descriptions of duties
- ◆ Test environment specifications
- ◆ Disaster and continuity plans
- ◆ Other aspects of computer operations

Next, a comprehensive policy should require that changes to these items not occur without proper approval and without documentation to reflect actual changes that are implemented. The policy should detail the change management process: request, review, approval, documentation, testing, implementation, and reporting.

The review, or approval process, must realize the necessity for levels of authority. For example, if a systems administrator needs to apply a critical patch, he should have the authority to do so. This does not mean that a blanket application of patches without testing should be allowed.

NOTE

Data Remanance *Data remanance* is the data that remains after data has been erased from physical media. Common misconceptions about deletion programs often leave quite a bit of data on the disk. First, PC delete programs merely remove the directory pointer to the data. The data actually remains. A low-level disk editor, a common utility, can be used to recover the data. Some disk wiping utilities do so by overwriting the data. Even this process might not remove all the data. Always look for a utility that overwrites data multiple times, or use some other method of removing data from the disk.

Deletion programs are not the only utilities to leave data chad on disks. Microsoft Windows Encrypting File System, if not patched, leaves bits of clear text data, called data shreds, when a clear text file is encrypted. This happens because a temporary file is created in the process and is originally not overwritten after the process is complete. Files encrypted from the start are not subject to this problem, and there are now free tools and a patch that can prevent the problem where this is not the case.

428 Part I EXAM PREPARATION

The approval for the application of a particular type of patch can vary. In some organizations this can mean exhaustive testing; in others a decision is made after review by a knowledgeable person. This process and policy involve the systems administrator.

Other types of modification, such as the implementation of new technology should be beyond the decision of a single systems administrator. This might require more stringent review that requires research and testing the impact on systems, network or application stability, cost, value, and product selection.

Regardless of the approval process, documentation must be changed to reflect current configuration and product mix. Documentation for related systems should also be reviewed. What impact does a new tape management system have on backup, offsite storage, recovery, collocation, compatibility, and training? Does new equipment bring new challenges in the availability of technical expertise, in application compatibility, or in the need for new auxiliary equipment and infrastructure? These questions should be answered prior to the change, but a review of related systems documentation and procedures can only occur once the product is installed.

Change management extends beyond documentation. If a new air conditioning system is to be installed, can it be scheduled for cooler months? Will the main power supplies need to be taken offline? (Should backup power be available and for how long?) If things don't work with the new system, can we fall back to the old?

By having a firm change management policy in place, the impact on the availability and stability of systems can be more reliably assured.

CASE STUDY: THE RUSSIAN HACK ATTACK

ESSENCE OF THE CASE

This case is an interesting one. The essence of the case involves the following:

- ▶ Strong encryption and other security measures did not prevent the hackers from success.
- ▶ Social engineering and common knowledge assisted the attack.
- ▶ A customer alerted the bank to the problem.
- ▶ Proper intrusion response practices stopped the attackers from being successful.

SCENARIO

In 1994 Vladimir Levin of St. Petersburg, Russia was able to hack into Citibank and steal \$12 million. He set up illicit funds transfers to banks in Israel, San Francisco, Finland, the Netherlands, Germany, Switzerland, and the Caribbean. He hired others to visit the banks to withdraw the money. How did he do this? The security at Citibank included strong encryption, solid procedures, and multiple-person control of transactions. Levin was able to determine the practices of Citibank and used this knowledge. Instead of directly attacking the computer systems, he spoofed real customer activity. He obtained account information and passwords and then ordered electronic transfers of funds from the customers' accounts to those he had set up.

ANALYSIS

This case shows how someone can use information about company operations to attack a company's assets.

This case of theft involved former employees of Levin's company, who moved to set up the bank accounts, which were used as repositories in the scam. In addition, they may have used the results of the prior successful attack on Citibank's computers by the Russian Hacker Megazoid.

Megazoid—a mathematical wizard, according to some accounts, or a group of hackers, according to others—may have provided information to Levin. Megazoid claims he remains anonymous for fear of criminal gangs anxious to acquire his skills. He claims he was able to navigate the Citibank network undetected for months.

continues

CASE STUDY: THE RUSSIAN HACK ATTACK

continued

He says he penetrated secret files, using a computer and modem he bought for \$10 and a bottle of vodka, as noted at <http://www.infowar.com/hacker/hacko.html-ssi>.

Official reports say that a large internal investigation cleared CitiBank employees of participation in the fraud. Bank security personnel in cooperation with the FBI were able to track illicit actions, arrest the moles, and gain information from them that eventually pointed to Levin and the company he worked for, AO Saturn. US authorities worked with the Russian Organized Crimes Squad. They then lured Levin to London where he was arrested. All but \$400,000 was recovered.

Levin was sentenced to three years imprisonment.

The service that Levin compromised was called the Financial Institutions Citibank Case Manager, which Citibank created in 1994 to allow

customers to transfer funds from their own accounts to accounts at other financial institutions around the world. To enter the system and transfer money customers were required to enter a user identification code and a password. Unlike similar operations by other banks of the period, Citibank did not also require a secure card for these transactions.

Think this is an isolated case? Think again. Security experts agree that it's not. They believe that banks hide information on successful hack attacks. They also believe that common penetration techniques will work equally as well at banks as they do in other industries.

For a peek into the techniques that might be used to do so, see the article, "How to Hack a Bank," at http://www.infowar.com/hacker/00/hack_052200a_j.shtml.

CHAPTER SUMMARY

KEY TERMS

- Administration or management controls
- Administrative management
- Annual Loss Expectancy (ALE)
- Application software maintenance controls
- Audit
- Audit and variance detection controls

Operations security involves figuring out what to protect, who to protect it from, who needs to have access, and what controls are available to help you protect it. Threats and countermeasures, auditing and intrusion detection, and OPSEC were discussed.

CHAPTER SUMMARY

- Auditing
- Banner grabbing
- Buffer overflow
- Capture
- Clipping level
- Controls
- Counteranalysis
- Countermeasures
- Corrective control
- Detective control
- Deterrent control
- Elevated privileges attack
- Ethical hacking
- Indicators
- Internet facing
- Initial Program Load (IPL)
- Intrusion detection
- Intrusion detection system (IDS)
- Intrusion prevention
- Intrusion prevention system (IPS)
- IPSec
- Operational controls
- OPSEC Process
- Packet
- Pen test
- Penetration testing
- Port redirection tool
- Port scanner
- Privileged instruction
- Privileges
- Promiscuous mode
- Protocol analyzers
- Qualitative risk analysis
- Quantitative risk analysis
- Ring zero
- Risk analysis
- Sniffers
- Switched networks
- Technical controls
- Threats
- Tip-off indicators
- Vulnerability
- War dialer

APPLY YOUR KNOWLEDGE

Exercises

7.1 Best Practices for Fax Services

Facsimile transmission at your company consists of several fax machines around the company. To send a fax someone must take a paper hard copy to the machine, load the document and punch in the recipients fax machine phone number. To receive a fax you must direct companies to use the fax number of the machine nearest you and retrieve the fax yourself. Fax machines are unmonitored and in public rooms.

1. Explain why the proposed controls listed in the worksheet outlined in Table 7.6 should be added to improve the security posture of your company's facsimile management.
2. Use your knowledge of operations security to mark your choices by placing an *X* in the Select column. Then describe why you made this choice in the third column.

NOTE

Fax Servers Rule! Fax servers are rapidly replacing individual fax systems. Fax servers can redirect received faxes to ordinary network printers throughout the organization. Fax servers can also direct faxes to the individual desktop and allow users to send faxes from the desktop—no hard copy is necessary. Ordinary scanners can produce electronic copies of paper documents that must be faxed. Some fax servers also allow the Print feature to be disabled.

TABLE 7.6

FAX CONTROL WORKSHEET

<i>Select</i>	<i>Control</i>	<i>Reasons for Choosing or Not Choosing Control</i>
	Require electronic receipt, no printout to uncontrolled fax machine.	
	Require monitors for fax machines.	
	Disable the print feature.	
	Direct printing of received faxes to network printers.	
	Install a fax server.	
	Require login to receive/send fax.	
	Require encryption of sensitive fax.	
	Require fax server to encrypt all sensitive documents. A separate fax server is supplied for sensitive transmittals.	

APPLY YOUR KNOWLEDGE**Answer to Exercise**

Table 7.7 provides the solution to Exercise 7.1.

TABLE 7.7**FAX CONTROL WORKSHEET ANSWERS**

<i>Select</i>	<i>Control</i>	<i>Reasons for Choosing or Not Choosing Control</i>
X	Require electronic receipt, no printout to uncontrolled fax machine.	Allowing faxes to print to unattended machines means that sensitive documents are available for theft or reading by anyone who happens to walk by. In addition, documents can be inadvertently picked up by someone honestly picking up his fax.
	Require people stationed at fax locations to monitor receipts.	This would assure some confidentiality but is not the best solution.
X	Disable the print feature.	Users may choose to print faxes (which may be sensitive documents) to network printers. Sensitive documents can still be left lying in unattended areas.
	Direct printing of received faxes to network printers.	This is not valid for the same reason that it isn't ideal to disable the print feature. Users might choose to print the faxes, which could then be left in unattended areas. This isn't a good situation when dealing with sensitive issues.
X	Install a fax server.	This can solve many problems but needs additional controls.
X	Require login to receive/send fax.	Excellent! Only authorized personnel can send and receive. Also ensures that the fax gets to the right person, and only that person.
	Require encryption of sensitive fax.	What, by policy? Who will remember?
X	Require fax server to encrypt all sensitive documents. A separate fax server is supplied for sensitive transmittals.	Yes. A technical solution exists that can ensure that sensitive documents are encrypted (presuming correct configuration is made and maintained).

APPLY YOUR KNOWLEDGE**Review Questions**

1. Describe the OPSEC process.
 2. Why are controls necessary for computer operations? Give examples of two types of controls.
 3. An IDS is what type of control? Why?
 4. Many operations security practices are based on security principles. Name and define two of them.
 5. What role does auditing play in operations security?
 6. What information can be gained by analyzing a capture?
 7. Discuss the proper role for penetration testing techniques.
 8. Which type of risk analysis uses hard statistical data to support its recommendations for countermeasures to threats? Illustrate your answer with an example.
 9. List and describe countermeasures to fraud.
 10. How should media be protected?
2. Which two methods can be used to purge RAM?
 - A. Degaussing
 - B. Clearing
 - C. Destruction
 - D. Removal of power
 3. Which of the following combination of duties into one job would violate the principle of separation of duties?
 - A. Configure security and systems programmer
 - B. Use of security systems software and auditing
 - C. Testing applications software and software quality control
 - D. System configuration and system troubleshooting
 4. Vulnerabilities in one's own network can be discovered by which of the following?
 - A. Clearing
 - B. A pen test
 - C. Looking at data remanence
 - D. Degaussing
 5. A technique used in risk analysis is which of the following?
 - A. Footprinting
 - B. Enumerating the network
 - C. ALE
 - D. OPSEC
 - E. Annual Loss Expectancy

Exam Questions

1. Which control is NOT an administrative or management control?
 - A. Personnel screening
 - B. Contingency planning
 - C. Separation of duties
 - D. Rotation of duties

APPLY YOUR KNOWLEDGE

6. Countermeasures to employee-related threats are which of the following?
 - A. Block all unnecessary inbound and outbound ports.
 - B. Eliminate banners.
 - C. Apply patches.
 - D. Bonding.
 7. A risk associated with administrative management is which of the following?
 - A. Ignoring controls
 - B. Building near explosion hazards
 - C. Championing professional development
 - D. Providing security training
 8. Antiviral products have been around for many years, yet we still have outbreaks of viruses and worms. The two most probable reasons for this are which of the following?
 - A. Gullibility of users.
 - B. Antiviral programs are not kept updated.
 - C. Antiviral programs cannot cope with the sophisticated virus programs written today.
 - D. Today's operating systems are more vulnerable to virus attacks.
2. Controls are necessary for computer operations to ensure that security is not compromised. A good control is separation of duties. Separation of duties prevents one person from being able to subvert or defraud or compromise the system. For example, an applications programmer should not also be a software tester. He might add backdoors to programs that would allow an attacker to compromise the system. As tester he could overlook this problem. Another control is setting permission on files. This technical control keeps data available for only those who should have the ability to access it. See the section "Identifying Available Controls and Their Types" for more information.
 3. An IDS system is an example of an audit and variance detection control because it looks for things which do not match the norm, and things which go against what is allowed. It also alerts an administrator about unusual circumstances. See the section "Identifying Available Controls and Their Types" for more information.
 4. Two security principles are separation of duties and least privileges. Separation of duties means to keep one person from entirely controlling a process that might allow them to defraud the system. Least privilege means to only give people the privileges that they need. See the section "Identifying Available Controls and Their Types" for more information.
 5. The role of auditing in operations security is to provide an audit trail or a list of what has happened to enable administrators to detect possible attacks and to determine if security policies are being fulfilled. See the section "The Roles of Auditing and Monitoring" for more information.

Answers to Review Questions

1. The OPSEC process is the process of looking at your company as the attacker would, discovering the information that he is seeing that might allow him avenues for attack, and then developing countermeasures so that this information is not

APPLY YOUR KNOWLEDGE

6. The analysis of a capture can provide information that allows detection of an attack, identifies the intruder, and provides forensic information for later analysis and possible prosecution. See the section “Detecting Intrusions” for more information.
7. Penetration testing techniques can discover vulnerabilities in your network and in your systems. A company should use these techniques after they have applied security hardening to their systems. The goal of pen testing is to find things that have not been discovered before, to catch configuration mistakes, and to have early warning of potential vulnerabilities. See the section “Penetration Testing Techniques” for more information.
8. Quantitative risk analysis uses statistical data to support its recommendation for countermeasures. An example is the Annual Loss Expectancy (ALE), which multiplies the loss potential times the probability of the threat occurring. See the section “Risk Analysis” for more information.
9. Countermeasures to fraud include separation of duties (ensuring no one person can do all of a process that would allow them to steal from or defraud the company), rotations of duties (ensuring no one is always doing the same thing), and mandatory vacations (fraud is often discovered when an individual is away). See the section “Establishing Countermeasures for Employee-Related Risk Analysis” for more information.
10. Media, tapes, and disks should be protected by labeling them, controlling access, keeping storage and usage area temperature controlled and clean, controlling and recording their movement, keeping them out of direct sunlight, and allowing

them to acclimate before using them when they are brought from outside to inside the building. See the section “Protecting Sensitive Information and Media” for more information.

Answers to Exam Questions

1. **B.** Contingency planning is an operational control. See the section “Identifying Available Controls and Their Types” for more information.
2. **B, D.** Random access memory can be cleared and will be cleared when power is removed. See the section “Protecting Sensitive Information and Media” for more information.
3. **A.** Configuring security is an administrative task. If a programmer configures security he might set it to be lax and then write programs that will more easily compromise the system. See the section “Identifying Available Controls and Their Types” for more information.
4. **B.** Clearing and degaussing are techniques to remove or destroy data on media. Data remanance is the data that remains after erasure of data from the system. See the section “Protecting Sensitive Information and Media” for more information.
5. **C.** Annual Loss Expectancy. See the section “Risk Analysis” for more information.
6. **D.** Bonding is the practice of paying a third party to insure the actions of an employee. It often includes some sort of a background check by the bonding agency and insures the company against fraud committed by the employee. See the section “Establishing Countermeasures to Employee-Related Threats” for more information.

APPLY YOUR KNOWLEDGE

7. **A.** If administrators flaunt controls, they set examples for their staff. They also are a greater risk, because they might have elevated privileges or access to confidential data. See the section “The Role of Administrative Management” for more information.
8. **A, B.** Many viral and worm attacks would not succeed if not for users who open attachments, respond to requests, download games, and so forth. See the section “Understanding Antiviral Controls” for more information.

APPLY YOUR KNOWLEDGE**Suggested Readings and Resources**

1. Fuld, Leonard M. *The New Competitor Intelligence: The Complete Resource for Finding, Analyzing, and Using Information About Your Competitors*. John Wiley & Sons, Inc., 1994.
2. Lee, Thomas. *Microsoft Windows 2000 TCP/IP Protocols and Services Technical Reference*. Microsoft Press, 2000.
3. Limoncelli, Thomas A. and Christine Hogan. *The Practice of System and Network Administration*. Addison Wesley, 2002. (Chapter 17, "Data Centers," and Chapter 25, "Organizational Structure.")
4. London, Robert W. "Employment Policies and Practices." In *Computer Security Handbook, Third Edition*, edited by Arthur E. Hutt, Seymour Bosworth, and Douglas B. Hoyt. John Wiley & Sons, Inc., 1995.
5. Northcutt, Stephen. *Network Intrusion Detection, an Analyst's Handbook*. New Riders, 1999.
6. Scambray, Joel and Stuart McClure. *Hacking Exposed Windows 2000*. Osborne/McGraw Hill, 2001.
7. <http://insecure.org/nmap/nmap-fingerprinting-article.html> (OS detection).
8. <http://samspade.org>.
9. www.arin.net/whois (ARIN).
10. www.atstake.com/research/tools/nc11nt.zip (netcat).
11. www.axent.com (AXENT).
12. www.cfisac.org/resource/OPSEC%20Indicators.com (Central Florida Industrial Security Awareness Council).
13. www.cybersafe.com (CyberSafe).
14. www.eeye.com/html/research/tools/index.html (EEYE).
15. www.foundstone.com/rdlabs/tools.phy?category=scanner (Foundstone).
16. www.iana.org/assignments/port-numbers (IANA port numbers).
17. www.infowar.com/hacker/00/hack_052200a_j.shtml ("How to Hack a Bank").
18. www.insecure.org/nmap.
19. www.iss.net (ISS).
20. www.nswc.navy.mil/ISSEC/Docs/Ref/GeneralInfo/opsec_basics.html (U.S. Navy).
21. www.nv.doe.gov/opsec/default.asp (Department of Energy, Nevada Operations).
22. www.nwpsw.com (netwcan tools prot).
23. www.opsec.org/ (OPS, the OPSEC Professionals Society).
24. www.snort.org.
25. www.systemexperts.com/win2k ("IPSec Filter," by Eric Schultze).
26. www.tripwiresecurity.com (Tripwire).

OBJECTIVES

Document the natural and man-made events that need to be considered in making disaster recovery and business continuity plans.

- ▶ Before you can successfully plan continuity and recovery, you must know the nature of the events that might cause you to use your plans. A simple listing enables discussion on their impact, assessment or risk, damages, and the operations necessary.

Explain the difference between disaster recovery planning (DRP) and business continuity planning (BCP) and the importance of developing plans that include both.

- ▶ Discussions of disaster recovery planning and business continuity planning often seem to be talking about the same thing. They both talk about calamitous events and what a business needs to do if struck by one. They both address the needs of this department or that and where to find help. To someone not involved in the planning effort, this is often confusing and can appear to be duplication. This section examines the difference.



C H A P T E R

8

Business Continuity Planning and Disaster Recovery Planning

OBJECTIVES

Detail the business continuity planning process.

- Explain the process of business impact assessment.
 - Define the process of developing the scope of a business continuity plan, including organization analysis, resources, and legal and regulatory requirements.
 - Develop business recovery strategies, including planning for crisis management; arranging for cold, hot, warm, and mobile recovery sites; communicating with personnel and management; and developing emergency response and implementation plans.
- The first step in planning business continuity is to understand the scope of the problem. A sound business impact assessment details the possible effect of every potential disaster. Every event can be analyzed as to its probability and how current business operation strengths and weaknesses impact the result. The planning effort asks the questions: Will operations be affected? Which operations are affected? Where will problems occur? For how long? How much will it cost? Does the organization have legal or regulatory requirements to fulfill? What about obligations to its employees and customers?

Next, an organization must determine which processes are most critical to business survival. For these critical operations, the cost and methodology of recovery must be determined.

Detail the disaster recovery planning process, including recovery plan development, implementation, maintenance, and the restoration of business functions.

- Define the process of recovery plan development.

- Describe emergency response, including the development of emergency response teams and procedures. Include disaster recovery crisis management and communication plans.
- Explain the necessary components of reconstruction procedures, including reconstruction from backup, movement of files from offsite storage, and loading of software, software updates, and data.

- *Disaster* is the name we give to an event that so cripples a business that operations can't resume for some lengthy period. When the event occurs, its first stage is often one of emergency. Every disaster recovery plan should encompass plans for action at the time of the emergency. A crisis can't be managed, but the response to one can be managed. Appropriate procedures, communication plans, and training provide the means to do so.

After the crisis is contained, an organization's personnel might be stunned into inactivity or busied with reconstruction. Proper planning provides the facilities, offsite storage of backups, tested procedures, alternative resources, and trained personnel necessary for the effort.

Explain the need for, and development of, a backup strategy. Include information on determining what to back up, how often to back up, as well as the proper storage facility for backups.

- Backup is not just the purview of IT. Formulas, manual files, business rules and procedures, and the collective knowledge of the organization are important in its recovery. Knowing what to back up and when to back up is critical. Full recovery depends on the provision of an appropriate storage facility as well as the proper procedural processes.

OUTLINE

Introduction	444	Testing the Plan	464
		Maintaining the Plan	465
What Are the Disasters That Interrupt Business Operation?	445	Defining Disaster Recovery Planning	466
		Recovering Data Processing	467
		Determining Recovery Plan Scope	468
		Creating Antidisaster Procedures	468
		Listing Necessary Resources:	
		Process and Site Selection Criteria	469
		Emergency Response Procedures	470
		Creating Step-by-Step Instructions	471
		Recording Important Contact Numbers	472
		Restoring Data Processing	472
Quantifying the Difference Between DRP and BCP	448	Developing a Backup Strategy	472
		Backup Procedures and Policy	474
		Vital Records Program	477
		Hardware Backups	478
		Alternative Sites	478
Examining the Business Continuity Planning Process	450	Chapter Summary	483
Determining the Plan's Scope	451	Apply Your Knowledge	485
Business Impact Assessment	452		
Gathering and Charting Information	454		
Validating the Process	456		
Reporting	458		
The BIA Process	458		
Developing Operational Plans	458		
Getting Help	460		
Reviewing Insurance	460		
Planning for Insurance Claim Processing	461		
Providing Item Recovery Details	463		
Implementing the Plan	464		

STUDY STRATEGIES

- ▶ BCP and DRP are, simply put, just a way of ensuring that some man-made or natural disaster does not eliminate the organization. An excellent way to study this topic is to use the methodologies explained here to develop plans for an organization with which you are familiar. Even if your job does not demand this knowledge or ability, you will gain a greater appreciation of the process and a better understanding of this domain by putting your quest for knowledge into a practical objective.
- ▶ If you do not feel you have the information available, or you feel the scope is too broad, select a department within the organization, or start your efforts by developing such plans for your family.
- ▶ Another strategy is to develop a plan based on what you know about a recent crisis. Would your organization have survived if your offices were in the World Trade Center on 9/11? What if they had been located in Southern California during the Northridge earthquakes of 1994, or on the coast of Florida during hurricane Andrew in 1992 or hurricane George in 1998? Whatever your choice for the exercise, involve yourself in writing a plan; don't just memorize terminology or attempt to learn this topic via osmosis.

Chapter 8 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING 443

“The Business Continuity Planning (BCP) and Disaster Recovery Planning (DRP) domain addresses the preservation of the business in the face of major disruptions to normal business operations. BCP and DRP involve the preparation, testing, and updating of specific actions to protect critical business processes from the effect of major system and network failures.

Business Continuity Plans counteract interruptions to business activities and should be available to protect critical business processes from the effects of major failures or disasters. It deals with the natural and man-made events and the consequences if not dealt with promptly and effectively.

Business Impact Assessment determines the proportion of impact an individual business unit would sustain subsequent to a significant interruption of computing or telecommunication services. These impacts may be financial, in terms of monetary loss, or operational, in terms of inability to deliver.

Disaster Recovery Plans contain procedures for emergency response, extended backup operation, and post-disaster recovery should a computer installation experience a partial or total loss of computer resources and physical facilities. The primary objective of the Disaster Recovery Plan is to provide the capability to process mission-essential applications, in a degraded mode, and return to normal mode of operation within a reasonable amount of time.

The candidate will be expected to know the difference between business continuity planning and disaster recovery; business continuity planning in terms of project scope and planning, business impact analysis, recovery strategies, recovery plan development, and implementation. The candidate should understand disaster recovery in terms of recovery plan development, implementation, and restoration.”

—Common Book of Knowledge study guide

This chapter covers Domain 8, Business Continuity Planning and Disaster Recovery Planning, one of 10 domains of the Common Body of Knowledge (CBK) covered in the Certified Information Systems Security Professional (CISSP) examination. This domain is divided into several objectives for study.

NOTE

Is Two Months Enough Time? The Ernst and Young study was released in March 2002. For comments on the practicality of creating a business continuity plan in two months, see http://www.computerworld.com/storyba/0,4125,NAV47_ST069705,00.html.

NOTE

Interagency Contingency Planning Regulation This regulation mandates that financial institutions in the U.S. will have a disaster recovery plan. It was developed by the Financial Institutions Examination Council, a "...formal interagency body empowered to prescribe uniform principles, standards, and report forms for the federal examination of financial institutions by the Board of Governors of the Federal Reserve System (FRB), the Federal Deposit Insurance Corporation (FDIC), the National Credit Union Administration (NCUA), the Office of the Comptroller of the Currency (OCC), and the Office of Thrift Supervision (OTS)." For more information, visit <http://www.ffiec.gov/>.

INTRODUCTION

In the aftermath of the 9/11 attacks on the World Trade Center in New York City, many companies rushed to update business continuity plans to include the potential for terrorist attacks. Others had no plan to update. In fact, two months later an Ernst and Young study revealed that 53% of U.S. businesses surveyed still had no business continuity plan.

I find it startling, but not entirely incomprehensible, that this is the general situation among U.S. businesses. After all, businesses have survived without such plans for centuries. Writing such a plan is no guarantee of survival. That's true, too. Why do business have a plan, and what does it encompass? How is one prepared, tested, and maintained? Is disaster recovery planning the same thing as business continuity planning? In today's at-the-speed-of-the-Internet world, where data is mirrored and co-located and stand-by systems and fail-over clusters are the rule, is backing up or a recovery plan even necessary? Where do all the parts fit in, and who is responsible for them?

These are the questions you should be able to answer about this domain, and these are the topics covered in this chapter. By way of introduction, let's review the reasons for having a plan:

- ◆ Studies indicate that nearly half of the companies that lose data in a disaster never reopen, and 90% of them are out of business within two years.
- ◆ Although countries might differ, in the U.S., the law does not explicitly mandate such plans, but it does mandate protection of business records. The Foreign Corrupt Practices Act of 1977 includes a requirement that compels corporations to keep accurate records and to safeguard company assets, and IRS 91-59 makes management responsible for record retention.
- ◆ Some types of businesses might be required to have a plan. The U.S. Federal Financial Examination Council, which regulates U.S. financial institutions, mandates a working disaster recovery plan for all U.S. financial institutions.
- ◆ Those companies that violate the law are subject to civil and criminal prosecution.

Chapter 8 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING 445

- ◆ Shareholders and employees of companies can, and have, sued companies for gross negligence in the absence of plans.
- ◆ Insurance companies might require the existence of such plans.
- ◆ Business partners, especially those who share access to their data systems, might insist on reviewing business continuity and disaster recovery plans.

Without a doubt, business continuity plans are necessary. The nature of the plan, and how they are prepared and managed, is the scope of this domain. It is my pleasure to introduce it to you. This chapter covers the following:

- ◆ Defining business interruption events
- ◆ Explaining the difference between business continuity and disaster recovery
- ◆ Examining the business continuity planning process
- ◆ Defining disaster recovery planning
- ◆ Discussing backup strategies

WHAT ARE THE DISASTERS THAT INTERRUPT BUSINESS OPERATION?

Document the natural and man-made events that need to be considered in making these disaster recovery and business continuity plans.

Disaster. The word immediately brings to mind catastrophic events. An earthquake killing thousands, a tornado flattening a town, a bomb exploding in a school house. Not one, but two, airplanes crashing into New York's World Trade Center. Disaster. Emergency vehicles. Press helicopters. Cries for help. We rush in (or want to) unprepared, unthinking, wanting to help. Others rush out. Quick! Escape, get out, run!

Our first thoughts, if not the way we carry them out, are correct. People are the most important resources. Saving life is the first goal of any response to an emergency or business interruption event.

NOTE

When Is an Event a Business Interruption?

Although everyone would consider a fire to be a business interruption event, few would see a small, quickly extinguished paper fire in a wastebasket as an event worthy of including in your plan, or as one that would trigger its operation. Granted, the little wastebasket fire needs attention, but the interruption is minor and the cost miniscule. Fire is an event to develop plans for. In that plan, perhaps, will be information that will qualify at what point the plan comes into being and perhaps referral to other policies and procedures that dictate activity for minor events of this type.

More lives can be saved if a plan has been developed to meet any emergency. With a plan, calm preparedness can reign, and where there is calmness, more lives can be saved.

What about the peripheral and inanimate objects the fleeing masses leave behind? What if the calamitous event is not a life-threatening disaster but nevertheless threatens the normal operations of whatever businesses are involved? What events should be considered in a plan?

In a business, any event that can interrupt its normal operation, which can negatively impact its people or its facilities, requires the creation of plan of action to deal with it. To create such a plan, you must first determine which events can threaten a business's ability to continue, and then, at what level those events trigger the operation of the plan.

The first step, however, is to list the events. Instead of merely adopting a prepared list, each business should create its own list, and the list should be reinspected at least annually to keep it up-to-date. The following list is the result of one business's recent discussion at the beginning of its business continuity planning session.

- ◆ Natural Events Including Weather
 - Earthquake
 - Hurricane or Heavy Rain/Wind
 - Blizzard or Heavy Snow/Hail
 - Tornado
 - Volcanic Eruption
 - Draught
 - Flood
 - Mudslide
- ◆ Terrorism, Sabotage, and Acts of War
 - Bombing
 - Kidnapping
 - Mailing or Otherwise Intentionally Spreading Life-threatening Bacteria or Viruses

Chapter 8 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING 447

◆ Accidents Including Environmental Spills

- Explosion
- Fire
- Power and Other Utility Outages
- Broken Pipes
- Hazardous Material Spill
- Nuclear Disaster
- Collisions from Vehicles—Trains, Autos, Boats, Aircraft

◆ Miscellaneous Events

- Explosion
- Hardware, Software Failure
- Strike and Picket Line
- Employee Evacuation, Absence
- Testing Outage
- Human Error and Omission
- Disgruntled Employee
- Malicious Mischief
- Vandalism
- Riot

These events are not ranked in order of severity or probability of occurrence. These steps must be taken and should be specific for each business location, but they should not be a part of the initial listing of events. In the beginning, every potential chance event—no matter how seemingly impossible—should be listed and not filtered.

Although it is important to make the list without speculation over which events actually represent a risk to this business, a risk analysis should be completed. To do so, review data on the FEMA site, community records of natural disasters and crime rates, as well as company history.

NOTE**So, Which Disasters Pose a Risk for You?**

Determine these by reading “Understanding Your Risks, Identifying Hazards and Identifying Costs,” a document available from the Federal Emergency Management Agency (FEMA), which you can read more about at http://www.fema.gov/mit/planning_toc3.htm.

You should also spend time hypothetically designing scenarios in which the unnatural disaster (terrorism, disgruntled employee, hacking attack, and so on), as well as the natural disaster, could cause you problems. For example, exactly what could a disgruntled former network administrator do to your network? What could a determined clerical employee do? (What access to data do they have or did they have while on the job?) Many of your planning efforts will revolve around mitigating the threat of business interruption due to these possible events.

QUANTIFYING THE DIFFERENCE BETWEEN DRP AND BCP

Explain the difference between disaster recovery planning (DRP) and business continuity planning (BCP) and the importance of developing both types of plans.

DRP and BCP often seem to be talking about the same things. The difference, however, is this: Disaster recovery is the process of bringing back into production a critical business process that has been crippled or destroyed by some catastrophic event. *Disaster recovery planning* is the process of developing a plan to do so; *business continuity planning* seeks to minimize the impact of catastrophic events on critical business processes, get the processes up and operational should some event occur, and bring the company back to full recovery after the immediate crisis has passed. Disaster recovery's emphasis has traditionally been focused on data processing and getting the data center functional. Business continuity considers both technical and operational business processes. It strives to keep the business solvent by determining which business operations are the most critical to the survival of the business and focusing recovery efforts on those operations.

Both planning efforts are necessary. Disaster recovery typically represents the immediate, short-term fix for affected processes. Business continuity represents the big picture. One without the other can lead to business failure. A business continuity plan without disaster recovery is not a business continuity plan at all.

Chapter 8 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING 449

How can a business continue if flood, fire, or some other event has knocked out network services, destroyed the data center, or injured or killed a large part of the workforce? If a company has a disaster recovery plan but does no business continuity planning, it might recover the data, data center operations, people, and facilities and yet the business might cease to operate.

Still, though, many companies don't seem to realize this. Perhaps it's the historical development of the process. It's not a bad idea to note the history behind the concept of planning for disaster. Understanding the rationale behind the various planning efforts as well as the differences between the two can help you avoid reliance on one or another of the planning processes.

Modern business continuity planning grew out of a need to develop plans to deal with the potential disaster of malfunctioning, damaged, or destroyed mainframe systems. These original efforts, called *disaster recovery planning*, focused on the capability of computer operations to deal with and recover from some disasters. Businesses recognized their growing reliance on their data systems and became afraid of the results should these systems be damaged or destroyed. Perhaps employees could return and ordinary facilities could be restored, but expensive computer systems and the data they held could not be so easily replaced. Elaborate plans to resume operations at remote sites, including standby equipment and data backup operations, became a necessary requirement for every data center.

Amazingly, at first, no one considered other aspects of business operation, nor what would happen if data systems survived and were again operational but the business could not function due to damage to other areas of the facilities, loss of critical employees, or loss of the ability to perform manual processing. No one paid much attention to the impact of monies lost due to lost business during the recovery operation or reserving emergency locations for people to work in, or what the impact of losing key employees in the disaster might be. Although the original emphasis on data system recovery was due to the business loss their demise meant, this reason behind the function was lost and the focus became simply keeping the systems running. I suppose businesses reasoned that disasters had happened in the past and businesses dealt with them.

Such inconsistent planning could—and did—lead to situations where the data center was again operational but the business was not. You might call this a business version of the cruel joke, “The operation was a success, but the patient died.”

Perhaps the cause was the movement of computers outside the data center and the need to plan for the recovery of distributed systems. Perhaps it was examples of business disaster that had to do more with procedures than with computing systems. Perhaps businesses with good disaster recovery plans failed after an interruption event. Perhaps it was a dawning recognition that data systems alone do not make the business. Whatever the cause, business viability became the goal. Business continuity requires more than data center recovery. It requires immediate response to a crisis; interim operation plans; recovery of data, equipment, and personnel; and finally complete restoration to normal operation. Business continuity planning is the creation of plans that ensure the continued operation of the business after some extraordinary event. The plan it produces must consider both the technical (disaster recovery planning) and operational restoration (business resumption planning) components.

EXAMINING THE BUSINESS CONTINUITY PLANNING PROCESS

Detail the business continuity planning process.

- **Explain the process of business impact assessment.**
- **Define the process of developing the scope of a business continuity plan, including organization analysis and resource, legal, and regulatory requirements.**
- **Discuss business recovery strategies, including planning for crisis management; arranging for cold, hot, warm, and mobile recovery sites; communicating with personnel and management; and developing emergency response and implementation plans.**

To respond to a crisis and restore normal operations, a business continuity plan must be developed. Although many steps must be taken in its development, many sources agree that the two most important items necessary for its success are backup and management support. Without backup, of course, there is nothing to recover, and without management support and guidance, no plan can succeed. Management support aides in obtaining money for mitigation processes (contracts for hot sites, duplicated systems, insurance reviews, and so on); time for planning, testing, and training efforts; and the support of the planning effort across boundaries of department, division, and role. It is management that eventually must decide how much money can be spent, and it is management support that ensures participation in the process. Fortunately, part of the planning process documents the financial impact of business interruption, and this information can ensure management's commitment to the planning process as well as plan implementation.

The business continuity planning phases are

- ◆ Determine the scope of the plan
- ◆ Perform business impact analysis
- ◆ Develop operational plans for each business process
- ◆ Test plans
- ◆ Implement plans
- ◆ Maintain plans

The following sections discuss each of the planning phases.

Determining the Plan's Scope

The scope of the plan must be derived prior to any planning process. Will the plan enumerate activity for the entire worldwide operations of a corporation, or will it focus on a specific facility?

NOTE

Audit Your BRP FEMA provides a complete series of checklists that cover the development of business recovery plans. The checklists cover four broad areas: executive awareness and authority, plan development and documentation, management and recovery team assessment and evaluation for effectiveness, and management and recovery team assessment of readiness and plan management. Although the checklists are directed at those developing a plan, in my opinion, they are far better used as an audit review of a functioning plan. They are available at <http://www.fema.gov/ofm/brecov.htm>.

Is this plan required for some new adjunct to the business: a new department, operation, or division? Should the plan address only a particular business process? Is it concerned with facilities, computers, and people or just one of these? Should the plan address all potential disasters or limit its efforts to a particular type?

Although every organization needs a plan that encompasses its entire operations and considers all possible business interruption events, if the organization has never had a BCP, it probably should focus first on only some part of the organization or recovery from a particular type of event. Another approach is to divide organization-wide planning efforts into localized or departmentalized planning efforts. These plans, when complete, can then be combined into a master plan for the entire organization. The master plan can address infrastructure, support services, and other areas that can impact multiple business processes and cross traditional business boundaries.

Regardless, the plan should not just address issues of putting critical components of the business back into operation; the scope of the plan should also address the legal and statutory elements that are a result of the business interruption. Legal and statutory elements can be fines that will be imposed due to late filing or completion of projects, penalties for not implementing mandated services and functionality, or the like. An example might be the fulfillment of new patient information privacy regulations as defined by the Health Insurance Portability and Accountability Act of 1996 (HIPAA), which can result in heavy fines and jail time. HIPAA outlines strict new guidelines on how every organization that deals with patient data must protect the privacy of the individual.

Business Impact Assessment

Observing the impact of disasters on others might have prompted the planning effort. But an organization must complete its own business impact assessment. Brainstorming a list of potential events that can affect the organization and trying to imagine how a particular business process would operate without some key resource is a good start.

Chapter 8 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING 453

However, the key to developing an effective business continuity plan is not only understanding the emergencies you might be faced with or which business operations they can affect, but in understanding what level of operations is necessary to fulfill the goal of keeping the business going. Unlimited resources for recovery will never be available, nor should that be your goal. The goal of recovery is to get critical services up and running to ensure the continuation of the business. Doing that requires a deep understanding of what these critical services are and the financial impact of their interruption.

A business impact assessment (BIA) is the process by which a business's critical services are identified and a maximum tolerable downtime (MTD) for each is determined. The MTD, sometimes also known as the recovery time objective (RTO), is the timeframe within which the critical service must become operational to ensure the business will survive.

A useful approach is to attempt to determine what will happen if each process can't function for several time periods. What will happen after one day of loss, after two, after a week? It is useful to attach dollar amounts in revenue loss, interest expense, discounts, fines, and so on—in other words, the total dollars over time that business interruption exacts. For each possible event, the operations that might be affected can be listed and a total financial picture determined.

These totals are useful in creating an awareness of the need for business continuity planning, and the impact of the loss of a single process helps to support funding for both pre- and post-event mitigation and recovery activities. Dollar figures also help separate processes into critical and noncritical operations and rank them in order of importance. However, two more factors should be studied. First, the interrelatedness of processes should be evaluated. Interviews with operations personnel might not reveal the true importance of a process. Understanding that some critical operation relies on this minor one might move the low-importance process to critical operation status. Secondly, some processes can survive moderate time periods of no function at all, whereas there may be a time at which critical operations must be resumed or no amount of money invested in the recovery process will be sufficient. Time-sensitivity is therefore a consideration. Hours of downtime for a Web site might be more devastating than days in a more traditional business.

NOTE**Life—the Most Critical Business**

Process Planners should make no mistake: The business impact analysis should rank processes that concern the lives of people as the most critical operations of all. Consider, for example, those operations that keep life support functioning first.

A figure often quoted by some insurance companies indicates that an outage of more than a few hours or a day or two would put them so far behind that they would cease to be able to do business.

Not every event has equal impact on all operations. Part of the analysis should determine which processes each type of event might affect. Tornadoes are unlikely to occur in New York state, hurricanes are unlikely in California, and earthquakes are less likely in Wisconsin. Each of these, though, has the power to wipe out a data center or a business, and the recovery planning for each can have similar requirements.

Finally, care should be taken to understand built-in or engineered fault tolerance. Saying that a system will be shut down because of the event has more impact if this means the physical system is destroyed versus merely temporarily shut down due to a power outage. In the former, replacement systems and data restore is necessary; in the latter, power restoration or even mitigation via power backup can mean minor interruption. This is why surveys and management level interviews are important but direct knowledge and probing follow-up questions are necessary.

Gathering and Charting Information

Obviously, a lot of research must take place. The one-to-one interview is a good and traditional way to do this. Other processes include small group meetings, video conferences, management-only interviews, and surveys. One of the goals is to determine the monetary loss that occurs when the business process is interrupted. Loss can be calculated by considering the following:

- ◆ Revenue loss
- ◆ Sales loss
- ◆ Interest lost on float

- ◆ Penalties for late payments to vendors or lost discounts
- ◆ Contractual fines or penalties
- ◆ Unavailability of funds
- ◆ Cancelled orders due to late delivery

Not all loss can be easily calculated in monetary value, but it should also be considered. Lists of other ways operations would be affected should also be made. These might include items such as loss of customer service capability; loss of the ability to help internal customers; and loss of confidence by customers, shareholders, employees, and regulatory agencies.

No matter how you conduct your research, sample questions can be found by examining published surveys. Arthur Hutt, in the *Computer Security Handbook*, is one such resource. Although it is a disaster recovery questionnaire and asks questions about computer applications, you could extend it to cover any business process or simply begin by abstracting questions for use in beginning your interviews. Hutt's questionnaire inspired Table 8.1, which could be used to combine the results of your surveys. Operations are listed down the side of the table; the impact, including loss in dollars, ranges across the top. It is meant as a start, to which you might add your own questions or adjust the timeframes. An e-commerce version of this business impact analysis table, for example, might substitute hours or minutes of operation in place of days. After the initial data is gathered, you can determine the operations most critical for business survival (those which would mean the most monetary loss if not quickly resumed). Next, calculate the Maximum Total Downtime (MTD), the time for which a critical operation can be down before the business loses its capability to survive. To do so, total the monetary losses over time and compare them to the loss that would be too much for the business to bear.

NOTE**Is a BIA Necessary for e-Commerce?**

Because e-commerce site requirements are 100% uptime, the MTD for e-commerce can be represented as 0. After all, high-volume sites might find that even tiny amounts of time offline result in staggering losses.

Recognition of these factors ensures management support, and initial funding plans often include complete redundancy for these operations.

Many sites are co-located (complete up-to-date copies of the site exist at other locations and can be almost transparently switched to if the site goes down). In the face of such operations, is there a need to perform a business impact analysis? Yes. The BIA can identify business processes that rely on the e-commerce activities, or which provide support for it so that appropriate plans can be made for them. Without a BIA, other activities—perhaps less obvious than being able to connect to the site but equally as important to business survival—might be overlooked. Can you imagine, for example, the impact if the catalog-ordering site was co-located but the warehouse was not? If a hurricane flattens the main site and the warehouse, the Web site might be operational elsewhere, but product still wouldn't ship.

456 Part I EXAM PREPARATION

TABLE 8.1**BUSINESS IMPACT ANALYSIS SURVEY RESULTS**

<i>Days from Event/Business Operation-Related Computer Applications</i>	<i>If Lost: Impact on Business</i>	<i>\$ Loss in Sales and Revenues</i>	<i>\$ Cost in Lost Clients</i>
Day One			
Operation 1			
Operation 2			
Operation 3			
Day Three			
Operation 1			
Operation 2			
Operation 3			
Day Ten			
Operation 1			
Operation 2			
Operation 3			
One Month			
Operation 1			
Operation 2			
Operation 3			

Validating the Process

To ensure more accurate findings, the business unit responsible for the process should validate the MTD derived from the information provided. An incorrect MTD can mean misdirected funds and resources. The MTD is used during the planning process to evaluate the recovery cost of making a system operational at certain times against the loss of revenue by its delay. Management's decisions then dictate the resources that can be made available for recovery. Business process owners need to consider which people need to be onsite and whether standby servers, alternative sites, or co-located services should be proposed.

Chapter 8 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING 457

<i>\$ Increased Expenses</i>	<i>\$ Additional Expenses to Restore Normal Business Operations</i>	<i>\$ Fines and Penalties</i>	<i>\$ from Legal, Civil Obligations</i>
------------------------------	---	-------------------------------	---

A 100% uptime (an MTD of 0), for example, can be met with alternative processing plans—for instance, hot sites, which have standby servers that can immediately take over operations, duplication of services at alternative places, and so forth. A plan to support processes with MTDs of several hours or days might include cold sites (sites with power and other facilities but no computers or software), restoration from backup, or even temporary alternative processing. Senior management will be asked to support the proposed effort to meet the recovery timeframe.

The correctness of the MTD should be evaluated prior to plan development. It will be much harder to revise or obtain approval for recovery plans at a later time in the planning process.

Reporting

A final report, called “BIA Findings and Recommendations,” is prepared. It should include an assessment of threats and vulnerabilities to time-critical business functions, document the impact (both operational and financial) on the business, and suggest a recovery approach that includes next-step recommendations.

This report should be circulated for final validation prior to publication. The results are often communicated to service organizations such as IT, network management, telecommunications, human resources, and the facility that supports each business unit. MTDs are often used during the rest of the planning process to determine measure, test, and deploy recovery processes.

REVIEW BREAK

The BIA Process

To summarize: The BIA process is a series of steps:

- ◆ Identify time-critical business processes.
- ◆ Identify supporting resources (personnel, facilities, technology, computers, software, networks, equipment, vital records, data, and so on) for the critical processes.
- ◆ Determine MTDs.
- ◆ Return to business units for validation.
- ◆ Provide the final report, including MTDs and recommendations for next steps, to senior management.

The results of this process are used to develop operational plans.

Developing Operational Plans

After validation and management approval of the BIA, a plan must be developed to ensure that critical operations will be available after a business interruption event. Each operation must be examined to determine which resources should be available.

Chapter 8 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING 459

The planning process can be divided into four phases:

- ◆ **Preventative measures**—Those operations that might prevent events, such as fire, or mitigate the effect of an event should it occur. Typical items in this part of the plan include fire and safety inspections, installation of fire detection and suppressant equipment, insurance review, attention to normal maintenance of equipment, data backups (including duplication of documentation, maintenance of backups, and storage of software offsite), training for employees, blast walls, and evacuation drills. They can also encompass a review of insurance for adequacy as well as training in the steps to be taken to ensure compliance with insurance policy requirements.
- ◆ **Emergency response**—Includes the actions taken immediately to avoid injury and loss of life, alert authorities, notify management, prevent additional damage, and (where possible) rescue critical data and equipment.
- ◆ **Recovery**—The process of putting critical operations back into operation. More information is available in the “Defining Disaster Recovery Planning” section later in this chapter.
- ◆ **Return to normal operations**—Transitional activity that returns the business to normal operations. This can include facility repair or replacement, establishment of new data and voice connections to support the entire operation, recall of employees, and the return of all operations to normal levels.

Plans must be made for each phase and include the activities that must occur, who is responsible for them, and what resources are needed. Once again, the business process owners are key players in the development of the plans. Because the BIA plan has identified the critical operations and the timeframes for their recovery, the business process owners can best define what is necessary to meet those timeframes. They should be trained in the process of evaluating alternatives for recovery, documentation of the strategies, and selection of key personnel to carry out the plans.

Some specific details that address these areas of the plan are

- ◆ Getting help
- ◆ Reviewing insurance

NOTE

Learning from the Past Examining the impact of disasters on business often suggests activities or approaches that can help mitigate the impact of future events. In the Chicago floods of 1992, many basement-level data centers experienced water damage, prompting many companies to redesign facilities to locate data centers above basement level. Hurricanes and other damage to data centers located on exterior windows walls have also resulted in the movement of data centers to interior areas of the building.

NOTE

Emergency Control Centers For some recovery operations, it might be helpful to plan emergency control center locations. These centers, located both within and outside the facilities, should include plan information such as an inventory of people, equipment, documentation, supplies, hardware/software, vendors, critical applications, data processing reports, communications capabilities, and vital records. During a crisis they can serve as communication centers and regrouping and recovery staging areas.

NOTE

Should Records Be Duplicated or Simply Protected in Fireproof Vaults?

Some records can't be duplicated. Money, equity certificates, and other forms of legal tender can't be duplicated. Other items are too numerous; consider for example the millions of pages of archived contracts and legal documents insurance companies and financial services organizations have. Any plan should consider the types of records that must be maintained and develop appropriate protective management and possible recovery efforts for each type of document.

- ◆ Planning for insurance claim processing
- ◆ Providing item recovery details

Getting Help

Plans for getting help should include specific steps to be taken during each phase. Contact information and notification steps should be prominently located and include

- ◆ **Telephone numbers of restoration companies**—In many cases restoration companies should be contracted just as hot, warm, and cold site vendors should be. In the event of widespread flooding, for example, restoration companies will be busy.
- ◆ **Phone numbers for insurance vendors.**
- ◆ **Instructions on proper notification**—This should also include information on what is covered and the approval process that is necessary before restoration work can be performed.

Reviewing Insurance

The planning process should include a review of insurance coverage. The goal is to determine whether current insurance is adequate and ensure that the recovery plan includes information that will allow those engaged in the recovery effort to best interface with insurance representatives for the best possible outcome. The time to learn about insurance is not when it is necessary. Insurance can provide funds to assist during recovery and restoration. Without insurance coverage, the business might be doomed.

Some items that should be questioned when assessing insurance policies are as follows:

- ◆ The type of risk covered
- ◆ The type of property policy valuation
- ◆ The need for specific additional insurance

Two types of risk can be quantified in the policy. *Named perils* specifies that the cause of the loss must be enumerated. If the cause is not listed in the policy, no coverage exists. Alternatively, *all risks* specifies that all causes of loss that are not explicitly excluded in the policy are covered.

Property policy valuation concerns the basis of compensation for loss. The two types—*actual cash value (ACV)* and *replacement cost*—both attempt to determine the cost to repair or replace lost or damaged items with those of similar quality and type. Actual cash value, however, deducts the value of physical depreciation, whereas replacement cost does not.

Many policies do not include coverage for the types of losses some business can incur. These might include the cause of the loss or simply might not cover the additional costs a business interruption event can generate. Coverage might be available but have to be purchased at additional cost. Each business will have to determine whether the special coverage is appropriate. Some of these items are

- ◆ **Business interruption insurance**—Covers lost earning and continuing expenses during business shutdown time.
- ◆ **Boiler and machinery**—Covers damage, replacement, and repairs necessary due to explosion of a steam boiler, pipes, engines, or turbines and mechanical breakdown.
- ◆ **Valuable papers**—Covers loss due to their loss or damage.
- ◆ **Accounts receivable**—Covers loss due to inability to collect.

Planning for Insurance Claim Processing

Restoration plans typically include those activities that follow the disaster recovery phase. But actions taken during and immediately following the crises can affect the capability of the business to restore operations to normal. These actions are those that would affect the capability of the business to effectively manage its insurance claims. It might seem awkward to create your recovery plans to guarantee the best possible insurance settlement, but the fact is that most insurance plans require businesses to take appropriate steps during and after business interruption events. If these steps are not followed, the insurance settlement might be much more difficult to claim and actually be reduced.

Each business should review its insurance plans with the insurance company representatives to ensure business recovery plans include the appropriate steps. Generic steps, and those typically useful in obtaining insurance claims, are detailed here:

- ◆ **Notify insurance company of claim immediately**—Give any details that are known and ask for assistance.
- ◆ **Secure the area**—Is it safe to enter? What needs to be done to ensure continued safety?
- ◆ **Restore fire protection**—Automatic or specific action might have removed power to sprinklers and other fire protection devices or otherwise removed any fire protection in place. If it is safe to do so, return operation of these devices, plans, and so on to protective status.
- ◆ **Prevent further damage/take action to minimize loss**—Perhaps water can be pumped out. Remove nondamaged goods to a place of safety and protection. If this is not possible, at least separate damaged materials, but do not destroy or trash them. Cover broken windows and holes in roofing as soon as possible. If possible, obtain emergency heat and dehumidification.
- ◆ **Provide security**—Guards might need to be posted or locks applied and barriers raised to keep out the press, the public, and employees not involved in damage assessment.
- ◆ **Take pictures and video of the site and damaged and undamaged property**—Documentation not only serves as a record for insurance claim purposes, but also can serve as a deterrent to theft.
- ◆ **Determine the cost of these and other temporary measures deemed necessary to resume operations and maintain security**—Often insurance can cover these costs and even provide emergency funding for these efforts. (You should, of course, be aware of these possibilities before an emergency.)
- ◆ **Obtain property replacement and repair costs from several sources**—Use internal engineering, operations, and maintenance personnel as well as outside contractors. Be sure to document the scope of activity this requires. Part of this process is determining what can be salvaged and repaired and what must be replaced.

- ◆ **Require all recovery personnel, including contractors, to log all activities**—Maintain a composite log.
- ◆ **Some steps are considered emergency response and simply must be done immediately**—Others, however, require approval. Determine the difference between the two and make sure that claims are submitted to the insurance adjusters appropriately and that a request for authorization to proceed is obtained. Although estimates allow beginning negotiations with restoration contractors, you will need to consult with the insurance adjusters before awarding contractors.
- ◆ **Partial payment might allow you to proceed with certain efforts.**
- ◆ **You might need to negotiate the final claim settlement.**
- ◆ **After the claim settlement is received, implement planning, acquisition, and installation of facility and resources.**

Providing Item Recovery Details

Although many recovery efforts need to be carried out by specialists in the field, knowledge of the commonly accepted practices of these experts should be common knowledge for internal recovery teams. Knowing these steps can either help avoid missteps that will prevent maximum recovery or provide staff with appropriate steps should recovery company representatives be unable to quickly arrive. Quick action, often that which can be done before the recovery company agents arrive on the scene, can be important. Think of it as first aid for critical data. An example of this is knowledge of the steps to recover from water damage to paper and tapes and disks. Please be sure to include in your plans the information given in this section, as well as the most current recommendations.

Quick action after a natural flood or water damage due to putting out a fire can salvage much paperwork. Specifically, water should be pumped out quickly and areas vented to allow air circulation. Cool temperatures can help preserve water-soaked documents, so storing documents in refrigerated trailers at 0° will help slow mold damage. Freeze-drying is also an effective technique. Before freezing, paper should be cleaned of debris and handled carefully. Paper that has coalesced into blocks should be kept in blocks and not pulled apart, and dehumidifiers can be used to dry documents. In addition, sterilization and application of fungicidal buffers can help prevent mold growth.

Computer tapes and disks need to be restored within 72–96 hours. Disks should be opened and dried with isopropyl alcohol and then placed in new jackets. Data can then be transferred onto new disks. Tapes should be freeze-dried or machine-dried with specialized machinery. Soot and smoke damaged disks need to be cleaned by hand and then data transferred to new disks.

NOTE

Don't Be a Statistic A 1998 Ernst and Young study found that only 27% of businesses with business continuity plans in place bother training staff in their operations.

Implementing the Plan

Implementing the plan consists of two phases. The first requires the acquisition of alternative equipment and locations, the acquisition of contractual arrangement with restoration specialists, and training of employees in their responsibilities and action during and shortly after each type of business interruption event. The second is the actual operation of the plan when an event occurs.

Testing the Plan

How do you test plan effectiveness? In the past, disaster recovery plans were often judged by a pass/fail on a computer recovery test. At other times, evidence of backup sites and redundant telecommunications were considered adequate tests of plan effectiveness. Neither of these is adequate, however. The plan must be exercised.

Several possible ways to test a plan are

- ◆ **Desk checking**—Reading through the plan and thinking how it would be used
- ◆ **Reviewing the plan for currency**—Examining the plan in light of new business processes, procedures, equipment, and interruption events
- ◆ **Performing full parallel system tests**—Testing backup equipment, software, data copies, and personnel at a hot site or alternative location
- ◆ **Running through scenarios and mock emergencies**—Having people respond by walking through their responsibilities as if it were a real emergency

Chapter 8 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING 465

- ◆ **Testing calls to contractors**—Finding out whether emergency personnel, facilities, and restoration specialists can be reached at any time of the day or night
- ◆ **Remote operations testing**—Moving employees to alternative sites and asking to operate remotely
- ◆ **Switching to the mirror system or site**—Performing a fail-over to a data vault
- ◆ **Reviewing insurance**—Making sure coverage is up-to-date and team members are aware of the steps to follow to ensure the best result
- ◆ **Testing by departments or business process groups**

Many organizations use a combination of the previous testing steps. A plan is considered valid and effective if it passes the following tests:

- ◆ Response is within the allowed time frame.
- ◆ Operations at alternative systems and locations are adequate.
- ◆ Backups can be successfully restored.
- ◆ Emergency personnel, service personnel, and contractors can be reached any time of the day or night.
- ◆ Team members are aware of the specifics of the current plan.
- ◆ Team members are able to perform associated duties.
- ◆ The plan is up-to-date.

Maintaining the Plan

No plan will stand the test of time. Processing routines change, hardware and software changes, employees come and go, and the number and type of business interruption events change as well. All these things and more require your planning efforts to be iterative. The business continuity plan must be reviewed on at least an annual basis. And, even more importantly, an examination of the relative portion of the plan should be made every time the business makes changes in its business processes. Change management should therefore include a review of the BCP as part of its checklist.

NOTE

Testing Insurance Any test of BCP should include a review of insurance. An examination of the policy should include a review of the adequacy of insurance coverage for recovery and restoration. Is there a need for, or is there adequate coverage for, vital records, equipment, restoration of data, and facilities? Are new coverage options available? Are some options no longer available?

A full review of the plan requires that each business process be examined to see whether the plan adequately addresses the needs of the current systems, equipment, facilities, and people. Among the items to review include

- ◆ Is the insurance plan up-to-date?
- ◆ Have new processes and equipment been added, and are they covered in the plan?
- ◆ Has team membership been adjusted to include or exclude changes in personnel?
- ◆ Is testing being done?
- ◆ Are there new types of events or changes in the likelihood of them occurring?
- ◆ Have mergers, acquisitions, or divestitures occurred, and has the plan been adjusted?

DEFINING DISASTER RECOVERY PLANNING

Detail the disaster recovery planning process, including recovery plan development, implementation, maintenance, and the restoration of business functions.

- **Define the process of recovery plan development.**
- **Describe emergency responses, including the development of emergency response teams and procedures. Include disaster recovery crisis management and communication plans.**
- **Explain the necessary components of reconstruction procedures, including reconstruction from backup; movement of files from offsite storage; and loading of software, software updates, and data.**

Because disaster recovery planning can be seen as part of business continuity planning, a similar planning process can be applied. The difference is that disaster recovery concerns itself with recovering or reestablishing technical operations of a particular process.

A good example—the one that, for many years, was the only planned recovery operation—is the recovery of data processing operations. Understanding and reviewing such plans allows you to adopt the planning process for the recovery of other technologies.

For the purposes of reducing redundancy, I will presume that the scope of your business continuity planning process encompasses disaster recovery planning and that a business impact analysis has been completed as part of that process, thus identifying the more critical applications and processes that are part of data processing. I will also assume that testing and maintenance portions of the disaster recovery plans can use the same instructions. Therefore, this section concentrates on the actual plan for post-interruption event recovery and the restoration of normal processing. Developing a backup strategy, a precursor to recovery, is detailed in a separate section.

Recovering Data Processing

The planning process for disaster recovery should include seven things:

- ◆ **The scope of the plan**—Including what is to be recovered and whether it's servers, data, or facilities.
- ◆ **Procedures that help to prevent disasters.**
- ◆ **A list of resources that need to be available**—Including an alternative site, equipment, data backups, personnel, and so on.
- ◆ **The backup strategy**—This ensures current data is available for restoration.
- ◆ **A to-do list for the emergency response process.**
- ◆ **Step-by-step instructions for implementing the plan**—This includes getting processes into operation.
- ◆ **Phone numbers of restoration and alternative sites**—Including business, home, off-hour numbers, cell, and other alternative numbers for locating your contacts at these companies.

Each of these is discussed in the following sections.

Determining Recovery Plan Scope

Just as the business continuity plan must first be scoped before planning can take place, the disaster recovery plan must identify which processes and equipment will be covered. The business impact assessment identifies critical data processing operations, and the disaster recovery planning effort determines exactly which equipment, software, facilities, environment, and personnel will be necessary to ensure their operation in the event of an emergency. For example, in evaluating alternative facilities, questions should be asked about equipment provisions and the need for climate control, security, raised floors, and so on. A distinction should be made, and special consideration taken for, data center-based operation versus distributed systems or Internet-based operation. If the planning process is the responsibility of IT, who will be responsible for systems that are critical but that are not the responsibility of IT? Many operations exist outside normal IT. An analysis of any company might find that financial operations, payroll, accounting, and even production systems are the responsibility of other departments. Additionally, data might be kept on user workstations or distributed to branch offices. Certain functions might be outsourced. You therefore have to ask which of these should be covered in the plan.

It is also wise to consider whether events such as mail storms, distributed denial-of-service (DDoS) attacks, and other types of attacks that cause business interruption are to be considered under the plan. No doubt, in many companies, the response to these events grew out of necessity and might not be formally codified as part of a disaster recovery plan. Should they be?

The answer often depends on the nature of IT within the organization. In a distributed infrastructure, control might be centralized or decentralized. Where centralized control is the rule, the scope of the plan should eventually encompass all IT operations; where decentralized control is the rule, plans will likely be created to only cover local IT operations. Having plans that fit the location and its needs is what's important.

Creating Antidisaster Procedures

Just as proper plant safety prevents accidents, safety and security procedures in the data center can prevent business interruption events or mitigate their impact. Although disaster recovery planning concentrates on dealing with emergencies, its study often supports the

NOTE**Disaster Recovery at Internet Speed**

When we think of disaster recovery, hot sites and the temporary movement of data processing to alternative sites are the first responses that come to mind. But those with e-commerce and other sites that have immediate technology needs, have long realized that a more immediate response to business interruption is necessary. e-Commerce sites can't wait for the activation of alternative data processing sites with reduced functionality. Other techniques and technologies must be considered. Typical responses include fail-over clusters, standby servers, co-location, and data vaulting.

Chapter 8 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING 469

practice of sound security procedures. During tests of the plan, improper procedures are often discovered. This information should feed back into procedural directives and employee training. Those doing plan reviews should also be trained to look for these issues.

For example, many companies purchase data safes—strong containers used to house onsite copies of backup tapes and offsite copies waiting for transport. Yet, these safes often remain open for the convenience of personnel. The safe, however, provides no protection unless it is sealed. Although procedures might even state this, employees need to be trained to shut the safe, and management needs to follow up to ensure the policy is being followed.

Other examples of anti-emergency procedures include the following:

- ◆ Locking hubs, routers, and switches in their own wiring closets instead of leaving them exposed in public areas or housed with public utility access points
- ◆ Limiting access to data centers, server rooms, and equipment closets
- ◆ Using approved fire-retardant materials in the construction of data centers
- ◆ Providing fire-extinguishing equipment and sprinkler systems where appropriate
- ◆ Performing background screening of employees
- ◆ Using antivirus products on gateways, servers, and desktops
- ◆ Using screening firewalls, routers, and so on at both egress and ingress points into networks

This list is not meant to be exhaustive; indeed, any good computer and operations security measure can be considered as lessening the chance of business interruption.

Listing Necessary Resources: Process and Site Selection Criteria

When critical business procedures must be relocated, or when backup includes parallel equipment, care must be taken to ensure a complete listing of resources is documented. Attention is usually focused on equipment needs, such as computers, wiring, and communications.

In addition, air-conditioning, fire-rated walls, dry sprinkler systems, fire abatement systems, equipment racking, power conditions, and UPS systems are necessary. There should also be plans for the movement of personnel and providing them with a place to work. Plus, the need for controlled access and security should be considered.

Considerations for site selection also require more than the capability to support processing. Sites should be evaluated to determine the capability of staff to get to them, their distance from the normal location, and their capability to manage any number of emergencies.

Emergency Response Procedures

Some business interruption events are more likely to cause panic than others. Who on your staff responds best in a crisis? Although it is difficult to anticipate how anyone will respond in an emergency, it is well known that people who are trained in the steps to take in an emergency do respond with more calm and are more likely to survive. Additionally, if other responsibilities need to be performed, where well-defined responsibilities are outlined, the outcome is more likely to be positive.

Therefore, you must create a list of instructions for all employees and train them in its use. Additionally, they must be empowered to act—there should be no question on the steps to follow. One thing must dominate all else: Life is the most important consideration. The first goal of any emergency response procedure should be to deliver people from life-threatening situations. The first step in any emergency response situation is to determine whether the situation is life-threatening. Although no one should attempt to set hard and fast rules on how to judge this, some training can be given to help supervisors and employees remain calm and make better decisions in an emergency. Giving examples of obvious threatening situations, such as a fire in the data center or an adjacent area, versus less threatening examples, such as fire in another building nearby, can help. The goal here is to keep people from blindly responding. No one wants staff to run out, leaving sensitive data exposed when there was time to secure it or transport it. On the other hand, no one wants to see employees die because they felt they had to shut down a server in an orderly fashion.

Procedures and training can help employees, supervisors, and managers judge when a specific response is required. Typically, separate procedures are warranted when lives are endangered. Clear instructions should indicate that when life is endangered authorities should be notified and an accounting of all people known to be in the building (employees, vendors, and guests) should follow evacuation.

A procedure is also needed for situations that are not life threatening. This list needs to be tested, and it should be updated periodically. Items on this list might include

- ◆ If programs are processing, shut down appropriately.
- ◆ Remove critical data files.
- ◆ Shut down equipment in proper sequence and shut off power.
- ◆ Establish damage control, such as covering equipment that can be exposed to water from sprinklers.
- ◆ If additional emergency control procedures exist, activate them if warranted.
- ◆ If appropriate, evacuate buildings.
- ◆ Reconvene at alternative sites.
- ◆ When appropriate, recall personnel for special assignments.

Creating Step-by-Step Instructions

You need to create step-by-step instructions on what to do if disaster strikes. These should include information on what to do, when to do it, and in which order to perform each step of the response and for each type of event. Copies of the instructions should also be kept offsite. More than one company has realized, too late, that plans were left back at the abandoned site. Employees should know where the plans are located and have practice in putting them into action.

Not all disaster recovery operations require movement to an alternative site. Instructions for these types of operations should be available as well.

NOTE

Flip Switch in Emergency A *shunt trip*, or emergency power shut-off switch, is often installed in a data center near an exit door. In case of fire, flipping the switch shuts off power, perhaps reducing the spread of fire and making the building safer for those fighting the fire. Please don't label the switch "Flip switch in emergency" with no additional information. An American company found out why the hard way: An employee became locked in the data center at off hours and pulled the switch, thinking it might provide means of escape. Well, he was rescued, but you can imagine the company's surprise when 150 Web servers suddenly shut down, removing the company's presence on the Internet.

Recording Important Contact Numbers

Not all the companies you work with will think to provide you with sufficient emergency numbers. It's funny, but for some reason, they tend to think of themselves as normal businesses and only provide daytime phone numbers for their personnel. Disasters, of course, are not considerate and often happen when businesses are closed. Take the time to have as part of your plan the additional off-hour phone numbers, and perhaps additional emergency numbers.

Restoring Data Processing

Plans for restoring normal operations after the emergency is over are often the purview of the business continuity plan. However, every disaster recovery plan should have procedures that indicate if this is so and who is in charge of the restoration process. Recovery plans can cover extended periods of time. Disaster recovery planning and business continuity planning need to detail procedures for operations over time. Plans that include movement to alternative sites should also have instruction for moving to other temporary facilities if that becomes necessary and detail the process for returning to the repaired or replaced permanent facility.

DEVELOPING A BACKUP STRATEGY

Explain the need for, and development of, a backup strategy. Include information on determining what to back up, how often to back up, as well as the proper storage facility for backups.

Backup is often defined as the placing of a copy of current data on tape media for storage. The goal is to have a snapshot of data from a certain point in time that can be used in an emergency to restore deleted, damaged, or otherwise missing data. A backup strategy, however, does not stop at providing the ability to recover data. This might be okay when equipment and facilities are not damaged or missing. However, a backup strategy includes the capability to move processing to alternative locations if necessary.

Every data center does a backup, don't they? IT audits still find sites for which backups are not done or for which they are not validated, carefully monitored and controlled, or tested. A data backup is insurance against the probability that something will damage data. Data, of course, can be damaged due to drive crashes or other media failure, accidental or malicious deletions, the introduction of bad data, or a virus or other attack.

There are many horror stories that recount failed data recovery because no backup existed or because the backup was not usable. Once again, the wise planner will assume the worst—all surprises will then be pleasant.

A comprehensive backup plan, including provisions for periodic testing, should be included in the disaster recovery plan. Backup plans include information on what should be backed up and when it should occur. Backup plans should exist as part of normal IT operation. Sometimes, however, a backup plan exists but is never implemented. There is no point in having a backup plan if you don't implement it. The plan should also include instructions on backing up data that does not electronically exist.

Many new technologies, such as mirrored systems, fail-over clusters, and data vaulting, provide alternatives to the simple restore and might cause some to question the necessity for backup. However, any system can fail, and a backup is always a cheap alternative to having no data at all.

The questions remain, "Is a sound backup policy in place? What is it? Is it used? Is it adequate? Is it tested? What are some generally agreed upon best practices? Is replacement, duplicate, or temporary use of hardware considered as part of the plan? Is movement to alternative sites arranged for?"

The planner should create plans based on current identification of critical systems, technology available, and recovery timeframe requirements. The wise plan includes the direct assistance of the technical individuals responsible for the systems in question. Items to consider are

- ◆ **Data backup**—Traditional copy to tape or other media.
- ◆ **Alternative sites**—Moving operations to other locations.

NOTE**Is Backup Always Necessary?**

Once, when I was teaching a class and we were discussing backup policies and procedures, I noticed that two ladies at the back of the room kept exchanging curious glances. I asked the class to explain their backup policies and procedures. After some discussion, I asked the ladies about their backup policies. "We don't back up," they said. The room sat in stunned silence. Astounded, I asked them who they worked for. "The U.S. government," they said. The room shook with laughter. It turned out, however, that the ladies had the correct backup policy for their environment. They managed a large database, and fresh data was downloaded every morning. No updating of the data was done at their site, and being without the data for the time it might take to download a new copy was an acceptable situation. In their case, it made sense not to back up.

- ◆ **Data vaulting**—Data, either the transaction or the data file, is transmitted to an alternative location in real-time. This can include the capability for a hot backup to immediately take over processing.
- ◆ **Co-location**—An exact copy, say of a Web or e-commerce site, is located at an alternative site or ISP. The co-located site is immediately ready to take over serving pages, accepting orders, and so on if a problem occurs at the main location.
- ◆ **Hardware backup**—Duplicate hardware is available either at the main site or alternative location, or both. It can immediately be put into service and the latest backup restored.
- ◆ **Hardware- or software-based redundant array of inexpensive disks (RAID)**—Fault-tolerant disk systems provide duplication of data or the capability to recover data in the face of drive failure. Several techniques are used. Data striping with parity provides on-the-fly recovery because the parity information enables data recovery should a single drive fail. Mirroring (two drives) and duplexing (two drives plus two disk controllers) write every bit of data twice. Should one drive fail, the other can take over.
- ◆ **Fail-over clustering**—Multiple processors operate in a cluster and provide the capability to automatically switch from malfunctioning units to functioning units.

Backup Procedures and Policy

Many companies adopt a policy of daily backups for servers, but the timing of backups should be a result of the amount of data that has changed and the critical nature of the data, as well as the capability of the system to back up when the data is not being used. Backups can be full or partial. In a *full* backup all normal files are copied. Exceptions to this are open files, database files, and some system files. Special backup agents can allow these files to be copied although they are online. Although full backups are preferred, *partial* backups provide a way for managing large amounts of data changes and large amounts of data. In many cases, the time to make a full or complete backup of all data can exceed the time allowed, especially if the data files to be backed up must be closed.

Chapter 8 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING 475

Partial backups can be made of data that has changed since the last backup.

Many companies adopt a strategy of making complete backups weekly, with partial backups made on the other days. In this scenario, a new, complete backup is made each week on a separate tape. Weekly tapes are kept for a month before being recycled, whereas daily partial backups must be kept for at least a week, depending on the type of partial backup made.

When a complete backup is made, each file backed up is marked as backed up. When a partial backup is made, however, only files that have changed are copied. Two types of partial backups exist. The *incremental backup* marks the copied files as being backed up, and subsequent incremental backups copy only files modified since the previous incremental backup. *Differential backups* also back up files that have changed since the last backup, but because these newly backed files aren't marked as being backed up, each subsequent backup also includes them.

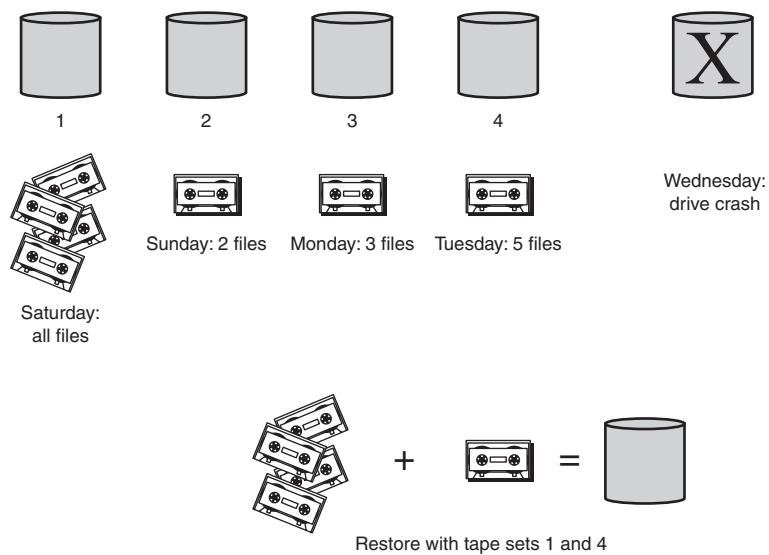
Examples of both incremental and differential backups are illustrated in Figures 8.1 and 8.2. In Figure 8.1, a complete backup is made on Saturday, which is then followed by differential backups during the week. On Sunday, two files, `productinfo1.dat` and `customerinfo2.dat`, are modified. The differential backup made on Sunday includes only these files and does not mark them as backed up. On Monday another file, `vendorinfo1.dat`, is changed. The Monday backup therefore includes `productinfo1.dat`, `customerinfo2.dat`, and `vendorinfo1.dat`. `pdata1.dat` and `pdata2.dat` are modified on Tuesday and included in Tuesday's backup along with the other three files.

Figure 8.2 shows the same systems, except this time an incremental backup is made on Sunday, Monday, and Tuesday. Incremental backups back up only files changed since the last backup but do mark the newly backup files as backed up. Sunday's backup contains the same files as that of Figure 8.2. Monday's backup, however, includes only `vendorinfo2.dat`, and Tuesday's backup includes only `pdata1.dat` and `pdata2.dat`. So as the week progresses, an incremental backup backs up less data each day than a differential backup, resulting in shorter backup times on consecutive days.

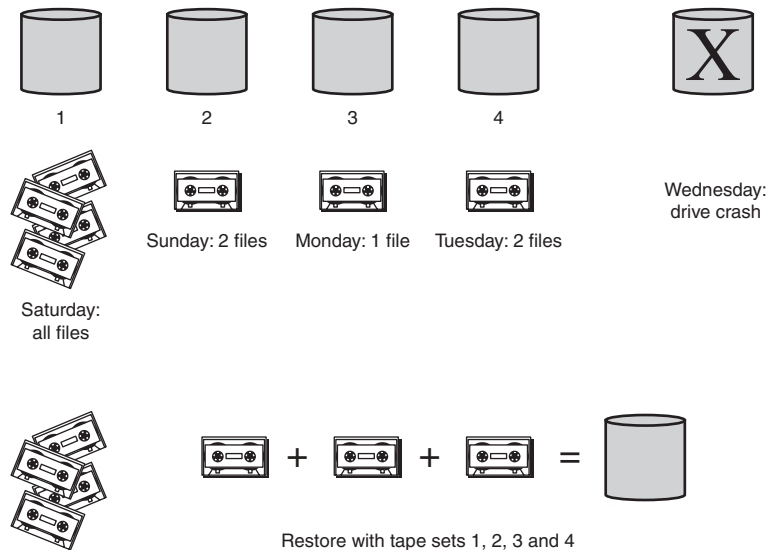
However, there's a bigger difference to keep in mind. If the hard disk crashes on Wednesday, the use of differential backups as in Figure 8.1 requires restoring only the complete backup made on Saturday and the partial backup made on Tuesday. In the Figure 8.2 scenario, all tapes are necessary—the complete backup and the partial backups from Sunday, Monday, and Tuesday. Planners and those responsible for backup and restore must understand the differences in tape sets necessary for recovery. Should Sunday or Monday's tape from Figure 8.2 be bad or missing, complete recovery is not possible.

Before you are tempted to require complete backups or differential backups, remember that if this includes huge amounts of data or data that frequently changes, there might be time and other constraints that require alternative backup procedures.

FIGURE 8.1
Full weekly backup with daily differential.



In many companies users are not allowed to store data on their desktop systems. This removes the issue of backups for desktops. But what about laptops and PDAs? What about desktop configurations? If users travel with their systems, they can't be expected to refrain from saving data on their machines. Backup systems such as Zip disks, read/write CD-ROMs, tiny hard drives, and other backup devices can be used as well as dial-up and Internet connections to store data. The company, however, must determine the procedures and policies that govern the backup of data stored on these devices.

**FIGURE 8.2**

Full weekly backup with daily incremental.

Another issue to consider is how and where tapes are stored. Both onsite and offsite storage should be arranged. Special cabinets and possibly special protective data safes might be provided.

Vital Records Program

As an addition to examining the critical business processes and the data systems by which they are supported, planners need to ensure the integrity and availability of vital records. *Vital records* are those that have critical importance to the organization and whose loss or damage would have a critical impact on business continuity.

Not all vital records are stored electronically, so provisions for securing them, such as duplicating microfiche and microfilm, paper, and other media, might be necessary.

In addition to onsite and offsite storage of current backups, many records must be archived for long periods of time to fulfill legal and regulatory requirements.

NOTE

How Do You Define Disaster? Today, authorities disagree on when business interruption becomes a disaster. A reasonable rule of thumb, though, is to consider an event a disaster when the entire facility is not functional and will not be so over a long period of time. This type of event usually means that processing will be moved to an alternative site. A catastrophe, on the other hand, includes major destruction of the facility and requires alternative facilities for possibly extended periods of time while new facilities can be built and equipped.

Hardware Backups

Data is not the only thing that might need to be recovered in the aftermath of some disaster. Hardware can be damaged, destroyed, or missing. A solid, current inventory of hardware will assist disaster recovery and restoration to normal processing. Depending on the critical nature of the processing, it might also be beneficial to maintain duplicate equipment; certainly, the availability of replacement equipment and the time it will take to do so weigh heavily in disaster recovery planning. Many interruptions will be localized, so it even makes sense to locate this duplicate equipment in the same building. Even a non-disaster (the result of system malfunction or failure) might need hardware to quickly resume service and prevent escalation into the disaster status. Of course, the cost of maintaining duplicate equipment should be factored into the decision to do so.

Alternative Sites

In picking alternative sites, many decisions must be made. Site type, location, size, and length of service must be determined.

Site type is usually defined as one of the following:

- ◆ **Hot**—Completely configured with equipment, systems software, and appropriate environment. It is only necessary to provide personnel, programs, and data, and recovery can be performed in hours. Usually reserved by paying a subscription cost, with additional charges for activation and daily use. Not intended for long-term use.
- ◆ **Warm**—Partially configured with the possibility of having peripheral equipment such as printers. Arrangements are made for this type of site if there is a good possibility of quickly acquiring replacement hardware. Might take days to make operational.
- ◆ **Cold**—Only the basic environment (wiring, power, air conditioning, and so on) is available. It can take weeks to make ready, so it is often used as a fall-back site from a hot site—in other words, a hot site is used while the cold site is being prepared.

Chapter 8 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING 479

- ◆ **Redundant**—It's set up exactly like the primary site.
- ◆ **Mobile**—A site configured in a trailer or van, it can be operational anywhere. It's often brought to the company to be used while the primary site is being repaired.
- ◆ **Hybrid**—It's some combination of these types of sites.

Information on alternative sites should be kept up-to-date. Constant contact and contract renewal should include the ability to maintain hardware and software compatibility. Imagine the surprise if you arrived at your hot site with data backups produced on a mid-range system only to find the site ready and waiting with a different system. Contracts should include the time during which the site will be available, what equipment is available, what if any staff assistance is available, when entry can be gained, and when tests can be conducted.

As an alternative to contracting with a specialized facility, some companies engage in mutual aide agreements. Each guarantees the other space, power, and possibly equipment to be used in an emergency. Each company should be as specific as possible, and contracts should be drawn up that specify what's available, when it's available, and for how long. Compatibility issues should also be addressed, and contracts should be regularly updated.

In addition to data, software and other information should be backed up. This might include

- ◆ Operating system software
- ◆ Programming languages
- ◆ Utilities
- ◆ Database management software
- ◆ Input, output documents
- ◆ Transaction logs
- ◆ System and audit logs

Several backup locations are usually used. The reason for multiple sites is that several types of problems might require the use of backups to restore systems. Many times a hardware failure requires the restoration of data. In that case, there is obviously no need to move to an alternative location and the data should be restored as quickly as possible.

Backups need to be close by. However, if the facility is destroyed, backups kept offsite will be available, whereas those stored near the data center might be destroyed. Some disasters affect several blocks or even entire cities or regions. In these cases, nearby offsite backup locations might also be destroyed. Having multiple backup locations ensures survival of data. Finally, some data needs to be kept for very long periods of time, so distant, more heavily protected repositories are desirable. Typical locations for backups include the following:

- ◆ A fire-resistant safe close to the computer room where most recent backups reside until transported to offsite storage.
- ◆ A fire-resistant vault in another building within a half-mile radius of the primary site. Backups can be stored here until they can be moved to a more distant site. The typical time frame is weekly.
- ◆ A fire-resistant vault at least 5 miles from the primary site.
- ◆ Underground, fire-resistant, and earthquake-resistant storage at least 50 miles away. Here records can be kept for many years.

NOTE

Dynamic Data Storage *Hierarchical storage management (HSM)* is the capability of a system to dynamically and automatically manage the storage and retrieval of online data files. Files that are infrequently used are automatically moved to storage media. Support for HSM is usually an operating system function on mainframe systems, but it might also be available on other systems such as Windows 2000. Special hardware is also required. In the event of system instability or malfunction, some data might not be online and thus unaffected. HSM devices might have removed data from a system, and thus there might be less data to restore. However, HSM should never be considered as an alternative to backups. Backing up all data is also required.

It's not enough to back up data. You must also know where it is kept, when the backups were made, what type they are, and how to use them to restore data. Good backup plans include instruction and information on

- ◆ Where backups are kept
- ◆ Labeling schematics for backup tapes
- ◆ Frequency of backup cycles and retention time
- ◆ Instructions on restoration, which include making a copy of the backup tape before attempting to use it in a restore
- ◆ How to recover from a failure during any step in the cycle
- ◆ Steps for special processing of special types of files, such as the agents necessary to back up databases online
- ◆ Documentation on backup files that create sets, such as transaction logs and database files

Chapter 8 BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING 481

- ◆ Locations of real-time or duplicate logs for transactions
- ◆ Information on ensuring the integrity of backup media
- ◆ The systems that require all files to be closed in order to be backed up and those that have available special agents that can be used in an online backup

Backup recommendations include

- ◆ Use a different tape for every day of the week.
- ◆ Create a weekly backup and use a separate tape for each week of the month.
- ◆ Verify each tape after creation.
- ◆ Check tapes for errors. Soft errors are recoverable; hard errors are not. A new backup on a new tape should be made.
- ◆ If unattended backups are made, make sure errors are logged to a file. Procedures should include steps for reviewing the log files.
- ◆ Clean the tapes.
- ◆ Use high-quality media.
- ◆ Change out tapes frequently, retire old tapes, and use new media.
- ◆ Label tapes immediately! Include the date of backup, the contents, and the machine backed up.
- ◆ Use a paper-based log to record when backups were made, what was backed up, and the location of the tapes.
- ◆ Test backups by doing a restore. Use the hot site if one is contracted.
- ◆ Log backup errors, exceptions, and anomalies.

NOTE

Is the Backup Good or Only the Header? Tape backup programs use different methods to verify the backup. Some check only the tape header; others confirm backup data is readable. If your backup program is only checking the headers, the backup could be unusable.

NOTE

Alternatives to Tape Tape has long been the media backup of choice. It's relatively cheap, widely available, and well understood. Its main detractors have been the time necessary to backup large amounts of data and the respective time to restore it. Alternative methods, such as parallel systems, fail-over clusters, and data vaulting, were developed to deal with time-critical applications.

As the cost of other electronic media, such as hard disk, CD-ROM, and DVD, continues to decline, businesses are considering and adopting these as the backup media of choice. Time for backup is reduced as is restore. In some cases, data can be considered to be online and instantly available. If these media are being used, backup procedures should be adjusted to work with them. Many of the same issues exist: Who is responsible for ensuring they are used? When are they used? Where are they stored? Care needs to be taken to ensure that appropriate copies are kept offsite so that recovery is possible should disaster require movement to alternative processing locations.

CASE STUDY: DOES BUSINESS CONTINUITY WORK?**ESSENCE OF THE CASE**

- ▶ A business continuity plan was in place; however, the unique way in which employees responded to a disaster ensured this company's continuation and subsequent successes.

SCENARIO

Yes (and the better your plans, the more likely it is). In the wake of the 9/11 attack on the World Trade Center, many businesses did not survive. But many did. The World Trade Center offices of bond trading giant Cantor Fitzgerald LP were destroyed, and 180 of its 733 employees were killed. However, Cantor was ready to trade two days later—in time for the September 13 reopening of U.S. Treasury markets.

According to an article in the December 13, 2001 issue of *Computerworld* (<http://www.cnn.com/2001/TECH/industry/12/13/redundancy.rebound.idg/index.html>) and information on the company's Web site (www.espeed.com), Cantor was able to do so because of built-in redundancy provided by its business-to-business online marketplace and IT services group, eSpeed (www.espeed.com), and because of the efforts of remaining eSpeed employees based in the U.S. and London. eSpeed had duplicated its IT services in a similar data center in the U.S. and was working toward uninterrupted uptime by linking both locations. Although that goal was not in place, each data center ran some of the services all the time, and periodic duplication of data from one to the other was ongoing. Additional backup facilities were provided by the London location.

Although the attack broke connections for U.S. customers, customers in Europe and Asia were unaffected. eSpeed also lost connections to banks, which meant it could not fulfill trade settlements.

CASE STUDY: DOES BUSINESS CONTINUITY WORK?

After the attack, employees worked around the clock to make sure the business could continue. They did so, they say, not because their jobs required it, but because they felt it was a way to reclaim what had been taken away. Nothing could restore the lives of those who died, but those who lived felt they could honor them by keeping the company going.

Shortly after the attack, trade settlement was outsourced to Automatic Data Processing (ADP). When the markets reopened, eSpeed was open for business and accepted the trades. Because bank reconnections were not completed by that time, however, it outsourced output to ADP for fulfillment.

Employees were successful. The company is doing well today and is caring for the families of the lost employees with health insurance and other benefits.

ANALYSIS

It would be nice to say that recovery was due to complete business continuity planning, but that was not the case. Outsourcing was not planned and practiced as part of a disaster recovery plan. Nevertheless, it was accomplished in just two days. Redundancy, dedicated employees, and the efforts of ADP made accomplishing the task possible. It almost seems—and the stories available for viewing on the eSpeed Web site verify—that the camaraderie and dedication of the employees was at least as important to the recovery efforts as the formal plan was.

CHAPTER SUMMARY

The business continuity planning and disaster recovery planning domain encompasses those activities required to ensure business survival in the face of events that interrupt its activities. Although the restoration of data processing and the recovery of computer operations are significant parts of that effort, technology recovery is not the entire story. Other business processes need to be evaluated, and their resumption planned, if a business is to survive. Business continuity planning might be best described as the merger of disaster recovery planning and business resumption planning.

KEY TERMS

- Business continuity planning (BCP)
- Business impact assessment (BIA)
- Business resumption planning
- Co-location
- Cold site
- Cooperative hot site
- Create and ship
- Data duplexing
- Data mirroring
- Data vaulting

CHAPTER SUMMARY

- Differential backup
- Disaster recovery planning (DRP)
- Fail-over cluster
- Federal Emergency Management Agency (FEMA)
- Full backup
- Full recovery test
- Hierarchical storage management (HSM)
- Hot site
- Hybrid site
- Incremental backup
- Maximum tolerable downtime (MTD)
- Mobile site
- Nonessential records
- Parallel test
- Partial backup
- Physical safeguards
- Procedural safeguards
- Recovery point objective (RPO)
- Recovery time objective (RTO)
- Redundant array of inexpensive disks (RAID)
- Redundant site
- Shunt trip
- Structured walkthrough test
- System downtime
- System outage
- Verify backup
- Vital records
- Warm site

APPLY YOUR KNOWLEDGE

Exercises

8.1 Researching Business Continuity Plans

The purpose of this exercise is to rate company plans for business continuity.

Estimated Time: 1 hour

1. Take the time to search online for companies or sites that provide information on business continuity or disaster recovery.

2. Rate these sites by analyzing the information they provide versus the marketing hype they offer. Which companies can provide evidence of their plans? Or, do the companies simply make promises? Create a chart, such as the one shown here, that includes your ratings. Evaluate the results.

<i>Site</i>	<i>Rating</i>	<i>Comments</i>
http://www.springboardhosting.com/products/managed_services/business.php?link=products	Just an ad; not much information	
http://www.disasterrecovery.com/	Contains a lot of information	A very good section on legislation and what is required as far as disaster recovery
http://www.riskconsult.com/home.html	Insurance/risk	Several articles on insurance, risk assessment
http://www.apexdm.com/	Contains just advertising	
www.tbicentral.com	Interesting articles	Must register

Review Questions

1. Where can you obtain information on the potential for specific natural disasters in your location?
2. Why should businesses have a business continuity plan?
3. Explain the difference between DRP and BCP.
4. Why should a business impact assessment be completed?
5. Identify the type of information you would collect from departments to determine whether a particular business process is a critical operation.
6. How do you determine the amount and nature of resources that will be prepared to successfully recover a business process?
7. Why is plan scope important?
8. If e-commerce operations are co-located, is a backup necessary?
9. Should the business recovery plan indicate anything that can be done before the interruption event occurs?
10. What's the difference between a disaster and a business interruption event?

APPLY YOUR KNOWLEDGE**Exam Questions**

1. A business impact assessment examines business processes to determine which of the following?
 - A. Which business processes are the most complex
 - B. Which business processes use computers
 - C. Which business processes are critical to the organization's survival
 - D. Whether a business process needs to be a part of the business continuity plan
2. A successful test of a business recovery plan has which following result?
 - A. A pass or fail
 - B. Demonstrated recovery of data from a backup
 - C. A visit to the hot site that reveals appropriate equipment is in place and operational
 - D. Information that can be used to make the plan more effective and knowledge of the readiness of the staff and availability of the equipment necessary
3. If a total disaster (the business facility is completely destroyed) occurs, which type of alternative site is best?
 - A. Hot site
 - B. Redundant site
 - C. Warm site
 - D. Cold site
4. Which requirement is most important during the analysis of the impact of business interruption on a particular business process?
 - A. How large the data file is
 - B. Current data duplication efforts already in place
 - C. The amount of money lost for every day of non-operation
 - D. Whether the operation directly impacts customers
5. The first step of any response to a business interruption event should be what?
 - A. If human life is at risk, evacuate the premises.
 - B. Call the proper authorities.
 - C. Secure critical or sensitive data.
 - D. Determine the source of the problem.
6. Business continuity planning is iterative. In which order should events occur?
 - A. Plan, train, test, revise
 - B. Plan, test, train, revise
 - C. Test, train, revise, plan
 - D. Plan, revise, test, train
7. Data management for e-commerce operations might include several functions designed to ensure 24/7 availability. If all of the following are being used, which of them can be eliminated without jeopardizing full data recovery in the event of a disaster?
 - A. HSM
 - B. RAID

APPLY YOUR KNOWLEDGE

- C. Daily backups
 - D. Data vaulting
 - E. Co-location
8. What is the first step in developing a comprehensive data management program?
- A. Ensure that all data systems are backed up.
 - B. Determine the location of all data.
 - C. Determine where critical data is stored.
 - D. Determine which data is most important.
9. You need to update a disaster recovery plan that was written when the only computers used in the company were mainframes. You are most likely to find that which of the following is true?
- A. Because processing is now distributed, a hot site is not necessary.
 - B. Because data vaulting is now practiced, data backup is no longer required.
 - C. Data might reside on user systems, and the plan must address responsibility for the backup of this data.
 - D. Individual departments have already developed comprehensive disaster recovery plans of their own.
10. What is the most important indicator of a successful business continuity plan?
- A. Strategies and operations are put into effect that prevent, reduce, or mitigate the impact of a disaster on the capability of a business to continue.
 - B. When tested, all operations such as data recovery, building evacuation, and location of alternative site personnel are successful.

- C. It covers every possible issue and resource necessary to recover operations.
- D. When a disaster occurs, people know what to do.

Answers to Review Questions

1. You can find historical information on natural disasters in your location by consulting old newspapers, historical associations, and municipal records. Information can also be found on the FEMA site (www.fema.gov). See the “What Are the Disasters That Interrupt Business?” section for more information.
2. Legal and statutory regulation of some industries might require a business continuity plan. Federal record keeping requirements also should be checked. See the Introduction for more information.
3. Disaster recovery planning is the process of creating a plan for the immediate recovery of technical business processes, such as those done by computer. Business Continuity Planning encompasses this, the mitigation of the effect of business interruption, the recovery of all operational business processes, and the restoration to normal function. See the section “Quantifying the Difference Between DRP and BCP” for more information.
4. A business impact assessment should be completed because it reveals the most critical business processes, allows their ranking, and produces a maximum tolerable downtime for each critical process. See the section “Business Impact Assessment” for more information.

APPLY YOUR KNOWLEDGE

5. A good indicator of whether a process is critical is if the business can survive very long without it. To find this out, you should ask what would happen if the process could not be completed for a certain time period (an hour? a day? a few minutes?); how much money would be lost, not earned, not collected, and so on; and what other processes would be affected. See the section “Gathering and Charting Information” for more information.
6. To determine which resources are necessary to recover a process, you have to look at the hardware, software, personnel, environment, and so on that the process is using today. Also important is knowledge of its reliance on other processes. See the section “Listing Necessary Resources: Process and Site Selection Criteria” for more information.
7. Plan scope is important for two reasons. First, if no plan exists, it is best to narrow the plan scope to more quickly and successfully create the plan. Often, choosing an area where disaster prevention procedures and mitigation can be established results in visible successes and enables future planning efforts. Second, management structure, corporate culture, or other political reasons might require some divisional development of plans. See the section “Determining Recovery Plan Scope” for more information.
8. Even though an e-commerce operation is co-located, a backup is necessary. Operational failure is not always so catastrophic as to require immediate change over to the alternative site. Problems can be as simple as an accidental deletion in an area of the site where the time to restore the data is minimal and can be tolerated. In addition, what happens if the alternative location is destroyed? See the section “Developing a Backup Strategy” for more information.
9. Business recovery planning includes a review of insurance, protective systems, and operational safety procedures to determine whether they are adequate. The planning group should always be searching for and recommending any additional items or procedure modifications that might prevent a business interruption or prevent it from becoming a catastrophe. See the section “Developing Operational Plans” for more information.
10. A business interruption event is any occurrence that halts normal business operations. A disaster is an event that cripples the organization so that the entire facility is not functional for a long period of time. See the “Hardware Backups” section for more information.

Answers to Exam Questions

1. **C.** Complexity is not a good indicator of the critical nature of a process. The simple process of checking picture badges against the person wearing them is critical to the security of the business. This process also does not use computers. Answer D might be an end result of the process but is not the best answer. See the section “Business Impact Assessment” for more information.

APPLY YOUR KNOWLEDGE

2. **D.** A simple pass or fail is difficult to determine because of the complex nature of the plan and the subjective nature of the process. Failure can be proven only if the business goes under, and that is impossible to determine in a test. Thus, determining what “passing” means is impossible. Recovering data from a backup only proves that the backup tape is good. Many other processes and events are required in most recovery efforts. Visiting the hot site can prove that equipment is ready—at that instant in time. However, each test of the plan teaches the business more about its operation and teaches the people who will need to perform the operations in the event of a real disaster. See the section “Testing the Plan” for more information.
3. **B.** The redundant site is exactly like the current facility, so it could more easily and quickly put the company back into operation. All the other alternative sites lack, or might lack, something that would mean a delay in resumption. (A hot site does not have your software loaded; a cold site does not have computers.) See the section “Determining Recovery Plan Scope” for more information.
4. **C.** The size of a data file can be important to consider in developing the procedure to deal with the operation, but it is not a good indicator of how critical the operation is. Existing data duplication is important because it means that less new expenditure will be required to provide adequate plans for its resumption. Customer-oriented applications are important and might actually be the most critical because they revolve around sales and the collection of money. However, some applications, such as customer support, less directly impact the bottom line.




The real indicator is the financial impact of the loss of the process. See the “Business Impact Assessment” section for more information.
5. **A.** Nothing is more important than human life. The absolute first response should be to prevent loss of life. If the risk is present, evacuate. See the section “What Are the Disasters That Interrupt Business Operation?” for more information.
6. **A.** Planning is necessary before testing. Training is the obvious second step. Testing reveals any need for revision. See the section “Implementing the Plan” for more information.
7. **A.** RAID provides fault tolerance. If one disk fails, data on the other disk(s) can be used immediately. Daily backups provide for restoration of data should other fault-tolerant methods fail. Data vaulting provides an additional copy of data at another location, and co-location provides a ready alternative processing site. However, HSM simply manages data, moving older data to less expensive storage mediums. It is not a good backup strategy because it does not represent additional copies of data and therefore can be removed without jeopardizing data recovery. See the section “Developing a Backup Strategy” for more information.
8. **B.** If you don’t know where all the data is, how can you manage it? Certainly backup is necessary, but what if you don’t know where all the data is? Knowing where critical data is located is important—do you know where all of it is? Knowing which data is most important is also vital—do you know where all of it is? See the section “Backup Procedures and Policy” for more information.

APPLY YOUR KNOWLEDGE

9. **C.** A hot site might still be necessary. Nothing in this description says the mainframe is gone, nor does it indicate that distributed systems might not be so critical that having an alternative, quickly available provisioned site might be important. Data vaulting is not a substitute for backup. Although departments might have plans, it is unlikely. It is, however, almost a surety that data resides throughout the company and determining where it is and how it can be backed up is now necessary. See the “Determining Recovery Plan Scope” section for more information.
10. **D.** It is not possible to ever know that all issues have been covered in a plan. Testing reveals whether those items tested work, but it does not prove the plan. Mitigation efforts are important but not as important as what people actually do when faced with a true disaster. See the section “Testing the Plan” for more information.

APPLY YOUR KNOWLEDGE**Suggested Readings and Resources**

1. Craig, Steven P. "Business Continuity in the Distributed Environment." In *Information Security Management Handbook, Fourth Edition, Volume I*, edited by Harold F. Tipton and Micki Krause. CRC Press, 1999.
2. Dorf, John, and Marty Johnson. "Restoration Component of Business Continuity Planning." In *Information Security Management Handbook, Fourth Edition*, edited by Harold F. Tipton and Micki Krause. CRC Press, 2000.
3. Hutt, Arthur. "Contingency Planning and Disaster Recovery." In *Computer Security Handbook, Third Edition*, edited by Arthur E. Hutt, Seymour Bosworth, and Douglas B. Hoyt. John Wiley & Sons, Inc., 1995.
4. Jackson, Carl B. "The Business Impact Assessment Process." In *Information Security Management Handbook, Fourth Edition, Volume 2*, edited by Harold F. Tipton and Micki Krause. CRC Press, 2000.
5. Jackson, Carl B. "Reengineering the Business Continuity Planning Process." In *Information Security Management Handbook, Fourth Edition*, edited by Harold F. Tipton and Micki Krause. CRC Press, 2000.
6. Peltier, Thomas R. *Information Security Policies and Procedures*. Auerbach, 1999.
7. Vallabhaneni, S. Rao. *CISSP Examination Textbooks, Volume 1: Theory*. SRV Professional Publications, 2000.
8. <http://www.brpa-chicago.org/BRPAinformation.html> (Business Resumption Planners, a nonprofit organization local to Chicago).
9. <http://www.disaster-resource.com/> (Annual Disaster Recovery Guide).
10. <http://www.drii.org/> (certified by the Disaster Recovery Institute).
11. <http://www.drii.org/lib/glossary.pdf> (glossary of terms).
12. <http://www.drj.com/> (Disaster Recovery Journal).
13. <http://www.drj.com/glossary/glossary.htm> (disaster recovery glossary).
14. <http://www.fema.gov> (Federal Emergency Management Agency [U.S.]).
15. <http://www.geocities.com/infosecpage/bcpdr.html> (business continuity and disaster recovery Web page of resources and free papers).
16. <http://www.globalcontinuity.com/> (portal for business risk and continuity planning).
17. <http://www.rothstein.com/data/index.htm> (catalog of disaster recovery books, tapes, CDs, reports, products, and so on).
18. <http://www.usfa.fema.gov/safety/sheets.htm> (generic safety sheets).



OBJECTIVES

This chapter covers Domain 9, Law, Investigation, and Ethics, one of 10 domains of the Common Body of Knowledge (CBK) covered in the Certified Information Systems Security Professional Examination. We have divided this domain into several objectives for study.

Explain the fundamentals of law.

- ▶ Without a proper introduction to the fundamental concepts of law, it will be difficult to understand the laws that can impact our use of computers, the resources we have to protect information systems, and the recourse we might have when our systems are abused by others.

Define what constitutes a computer crime and how such a crime is proven in court.

- ▶ You need to know the role of the law in computer security, especially criminal law. You should also learn what constitutes a computer crime and how such a crime is proven in court.

Explain the laws of evidence.

- ▶ Courts take action based on the establishment of facts based on evidence. The way in which a court deals with evidence depends on the laws or rules of evidence. By understanding those laws, a computer professional will better be able to design security systems and execute investigations about security incidents.

Introduce techniques for obtaining and preserving computer evidence.

- ▶ Use of the proper techniques for gathering evidence will enhance its value in civil trials and criminal prosecutions. This chapter introduces the principles that should guide the acquisition of evidence in any computer investigation.



CHAPTER 9

Law, Investigation, and Ethics

OBJECTIVES

Identify and plan for computer security incidents.

- ▶ A computer security professional should be capable of helping an organization prepare for computer security breaches. Preparation requires knowledge of the different ways in which someone might challenge computer system security and methods for responding to an incident when it occurs.

Discuss computer ethics.

- ▶ Does the law have anything to say about ethics and computers? What about the self-imposed rules of computer scientists and users of computing facilities? We take for granted that others have the same beliefs we do, but perhaps it's time to clearly state what that means. To start, you should investigate what has been said in the past about your ethical responsibility toward computing facilities and other computer users. Furthermore, to become a CISSP you must sign a statement of ethics. You should understand the statement you are signing and how it relates to the security professional's job.

OUTLINE

Introduction	496
Fundamentals of Law	497
Intellectual Property Law	498
Patents	498
Copyrights	498
Trade Secrets	499
Sale and Licensing	499
Privacy Law	500
Government Regulations	502
Criminal Law and Computer Crime	503
Computer Security Incidents	505
Advance Planning	506
Computer Crime Investigation	507
Legal Evidence	509
Credibility or Weight of Evidence	510
Proof of Authenticity	511
Hearsay	511
Best Evidence Rule	511
Chain of Evidence	512
The Fourth Amendment	513
Computer Forensics	513
Computer Ethics	517
Chapter Summary	521
Apply Your Knowledge	523

STUDY STRATEGIES

- ▶ The best way to learn the material in this chapter is to read it with an active mind. Don't just try to memorize it. Think about it. Notice the interrelationships between the different subjects. It is not entirely predictable what subjects might be covered by this portion of the CISSP exam. By getting a feel for what is right and wrong, you'll better be able to select the best answer on each exam question.
- ▶ This chapter guides you with questions to contemplate as you read. Thinking about the questions should help you remember the concepts.
- ▶ This chapter can't cover every fact of law, investigations, or ethics that might possibly be included on the CISSP exam. It is recommended that as you study sections in this chapter, you also read the additional reading and background material cited throughout the chapter and at the end of the chapter.

“The Law, Investigations, and Ethics domain addresses computer crime laws and regulations; the investigative measures and techniques which can be used to determine if a crime has been committed; methods to gather evidence if it has; as well as the ethical issues and code of conduct for the security professional.

Incident handling provides the ability to react quickly and efficiently to malicious technical threats or incidents.

The candidate will be expected to know the methods for determining whether a computer crime has been committed; the laws that would be applicable for the crime; laws prohibiting specific types of computer crime; methods to gather and preserve evidence of a computer crime; investigative methods and techniques; and ways in which RFC 1087 and the (ISC)² Code of Ethics can be applied to resolve ethical dilemmas.”

—Common Body of Knowledge study guide

INTRODUCTION

The topics of this chapter all interrelate. Computer crime laws are based on rules of ethics. The prosecution of a computer crime depends on the availability of evidence. And, evidence is gathered through investigations.

Often, breaches of computer security are also crimes for which perpetrators can be prosecuted in court. Gathering evidence for prosecution therefore might be one of the objectives in a response to a computer security incident. Some types of evidence are better than others, and often the difference depends on the techniques used to gather and preserve the evidence. This chapter shows the relationships between security breaches, law, incident response, and computer evidence forensics. It also introduces the ethical responsibilities of computer security professionals.

Except as otherwise indicated, the laws addressed in this chapter are American laws. You should also recognize that this chapter provides only a very general statement of the law, and nothing in this chapter is legal advice for a particular situation.

Before you proceed, ask yourself some questions. What should the law deem to be a computer crime? What must happen before the government can brand a person as a computer criminal? How should a court know whether any piece of computer evidence is what it appears to be and is not fabricated or altered? What procedures or ethical standards should a computer security professional follow so courts and other law enforcement authorities will believe what the professional has to say about any particular incident?

FUNDAMENTALS OF LAW

Explain the fundamentals of law.

Laws in the United States are either *federal*, which apply nationwide and originate from legislation enacted by the U.S. Congress, or *state*, which apply only within the borders of the state in question. Often, the subject matter covered by federal and state laws can overlap. For example, unauthorized intrusion into a bank's computers might violate both the federal and state computer crime laws. The intruder could be convicted under both federal law and state law, and the law enforcement authorities having jurisdiction over investigation and prosecution of the matter might be both federal and state.

Criminal laws authorize the government to punish wrongdoers with financial penalties and incarceration. To convict a suspect under criminal law, the government must meet a high standard of proof—*proof beyond a reasonable doubt*—that the suspect intentionally did something wrong.

Civil laws, on the other hand, enable private parties to enforce their rights—such as contract, tort, and property rights—through court orders and monetary awards for damages. An example of a tort is negligence, where one party injures another by failing to exercise ordinary care to avoid injury to the other. To win relief under a civil lawsuit, a plaintiff must satisfy a lower standard of proof—*proof by a preponderance of the evidence*—that she is entitled to relief.

Administrative law allows government agencies to interpret the laws they administer through official statements or regulations and to enforce those laws through investigations, fines, and other sanctions.

NOTE

Reasonable Doubt Criminal prosecution requires a higher standard of proof—proof beyond a reasonable doubt—that the suspect intentionally did something wrong.

Intellectual Property Law

Suppose an entrepreneur has an idea for a new technology. How would she protect rights to the idea? The major categories of intellectual property law available are

- ◆ Patents
- ◆ Copyrights
- ◆ Trade secrets

Pirates who violate these laws can be liable for civil damages to property owners and even be subject to criminal prosecution.

As you read about patents, copyrights, and trade secrets, notice that these intellectual property laws do not protect all the ideas an entrepreneur might devise.

Patents

A patent grants to its owner the exclusive right to make, use, or sell an invention covered by the patent. A patent can cover a physical invention or a business process, such as a unique process executed by software. To obtain a patent, an inventor must apply to the U.S. Patent and Trademark Office (USPTO). Often, the inventor must wait two or three years before the USPTO decides whether to grant the patent.

Copyrights

Copyright law grants to the owner of a copyright the exclusive right to copy and make derivative works from the copyrighted material. Copyright covers expressions of ideas, such as written words, pictures, sounds, software code, and even live performances. But copyright covers only the expressions of the ideas, not the ideas themselves. For example, if an entrepreneur has an idea for a scrumptious pizza recipe, and she writes that recipe in a book, she then owns the copyright to the words in the book (the expression), but she does not own a copyright to the combination of ingredients and techniques that are used to make the pizza (the idea). Copyright applies automatically to original material as it is created. Copyright law grants to copyright owners special advantages if they mark their material with copyright notices and register their material with the U.S. Copyright Office.

Intentional copyright infringements for commercial advantage or financial gain can be a crime. Also, the Digital Millennium Copyright Act (DMCA) makes it a crime to make, sell, or distribute products or services intended to circumvent the encryption or other technical devices that copyright owners use to protect their copyrighted material. It also makes it a crime to break encryption or other devices for the purpose of gaining unauthorized access to copyrighted material. Criminal prosecution under the DMCA requires that the perpetrator act for the purpose of commercial advantage or financial gain.

Trade Secrets

Trade secret law allows the owner of a trade secret to prevent others from using or exploiting the secret. A trade secret might be something like a customer list or an algorithm for searching through data on a network. Trade secret law applies automatically to information a company treats as a trade secret. (It does not apply to a pizza recipe published in a book because publication makes the recipe no longer a secret.) To maintain trade secret rights over information, companies must take steps to ensure the information does not become known to the public. Therefore, companies protect their secrets with security methods (encryption, logging copies, and so on) and by asking employees and business partners to enter agreements of nondisclosure. Theft of trade secrets can be a crime.

Sale and Licensing

When a programmer or a contractor is hired to write software, the employer typically obtains an agreement that all the programmer's or contractor's work product (inventions, copyrights, and trade secrets) are sold and assigned to the employer. This arrangement is known as *work for hire*.

But when a software developer creates software for the purpose of marketing it to multiple user customers, the developer typically grants to each customer only a *license*. A license is typically a contract that allows each customer to use the software (and the patents, copyrights, and trade secrets therein), under restricted terms, but does not allow the customer to remarket the software as its own. A license typically means a right to use but not to own.

Privacy Law

Do people have a general right to privacy of information about them? Another way to ask the question is this: When can a company be liable for violating someone's private information? As you learn the answer to those questions from the following material, observe the key role published privacy notices or policies play.

The United States has no comprehensive national law on privacy. U.S. privacy laws tend to apply on a sector-by-sector basis.

One such sector is healthcare. State laws and the federal Healthcare Insurance Portability and Accountability Act (HIPAA) generally require healthcare providers to maintain the confidentiality of patient information.

The federal Gramm-Leach-Bliley Financial Modernization Act requires financial institutions to give customers notice about how their private information will be protected or shared with third parties. Under the act, financial institutions are free to share information so long as they give customers notice and, in some cases, the opportunity to opt out of information sharing. Failure of an institution to abide by its notice can lead to liability.

The Privacy Act limits the ability of federal government agencies to disclose to the public or other agencies information they have about individual citizens.

NOTE

Privacy Policy Liability Failure of a company to abide by its published privacy policy can lead to liability.

Generally, no American law requires that companies post privacy policies with respect to people who visit their Web sites. However, many companies do elect to post privacy policies to make visitors feel more comfortable. Such a policy might say something to the effect that the company will not share with third parties private information collected from visitors. This policy is like a contract, and failure on the part of the company to comply with it can lead to civil liability. For example, US Bancorp paid a total of \$7.5 million to settle charges that it used private customer data in violation of a privacy policy it posted on its Web site. See http://www.ag.state.mn.us/consumer/Privacy/PR/pr_usbank_07011999.html, http://www.ag.state.mn.us/consumer/Privacy/PR/pr_usbank_06091999.html, and the September 1, 2000, press release at <http://www.firststar.com/about/ii-news-fr.html>.

Generally speaking, employees have no right to privacy when communicating through corporate information resources if the employees are informed in advance that they have no privacy. Therefore, many corporations publish notices to employees to the effect that management might monitor their email or other electronic communications.

These notices can be communicated to employees as agreements, to be signed by the employees, showing they acknowledge they have no privacy rights relative to data on company machines. Further, these agreements can include company security policies that employees are expected to follow. Requiring employees to sign written security policies is a practical way to persuade employees to protect company resources and information. Training and awareness programs that educate employees about security are also good techniques.

In contrast to the U.S., the European Union (EU) has more comprehensive rules on individual privacy. Traditionally, these rules have included restrictions on “transborder data flows” that would allow private data to flow to countries whose laws would not protect that data. The European Union’s Directive on Data Protection forbids the transfer of individually identifiable information to a country outside the EU unless the receiving country grants individuals adequate privacy protection.

To establish that data sent to the U.S. is granted adequate privacy protection, the EU and the U.S. government have negotiated a *safe harbor*. Under the safe harbor, participating U.S. companies voluntarily agree to protect personally identifiable information from the EU by, among other things, granting EU citizens the rights to the following:

- ◆ Notice about which data will be collected and how it will be used
- ◆ Choice about whether data will be collected
- ◆ Access to collected data
- ◆ Reasonable protections for accuracy, integrity, and security of collected data
- ◆ Rights to seek redress for abuse of data

See safe harbor materials at <http://www.export.gov/safeharbor/>. These rights are consistent with commonly recognized fair information practices.

Also relevant to the law of privacy in the U.S. is the Fourth Amendment to the U.S. Constitution. See the section “The Fourth Amendment,” later in this chapter.

Some companies employ privacy officers to monitor how private information is used within the organization and make recommendations to management for protecting privacy better. The presence of a privacy officer in a company is evidence to regulators and courts that the company is making a good effort to address the often difficult challenge of protecting privacy.

Government Regulations

Some specific laws mandate that enterprises institute information security controls.

The federal Foreign Corrupt Practices Act (FCPA) requires publicly owned companies to maintain adequate books and records and an adequate system of internal controls. Normally, the FCPA is enforced as administrative law by the U.S. Securities and Exchange Commission.

The federal Gramm-Leach-Bliley Financial Modernization Act, and official guidelines published under the act, require financial institutions to implement a security program to safeguard private customer information in their possession. See, for example, the guidelines published for banks by the Office of the Comptroller of the Currency at 12 Code of Federal Regulations Part 30, Appendix B.

To stem the transfer of military or strategic capabilities to undesirable countries, the U.S. Export Administration Regulations require that exporters obtain licenses before they export certain high-performance computers and microprocessors, as well as strong encryption. The U.S. Commerce Department’s Bureau of Export Administration (BXA) administers and enforces these export controls. Noncompliance can lead to administrative sanctions and criminal penalties. Accordingly, software containing cryptography functions commonly comes with a license that forbids the licensee from taking the software outside the United States.

CRIMINAL LAW AND COMPUTER CRIME

Define what constitutes a computer crime and how such a crime is prosecuted in court.

Criminal laws punish serious offenses against society. Under the criminal laws, the government, acting through a prosecutor who appears before a court, can convict a suspect such that he obtains a record as a criminal and can be subject to penalties such as monetary fines and incarceration.

When is it that the government should have the power to convict someone as a criminal, strip him of his liberty and lock him in prison for breaking into a computer system? All criminal convictions, whether computer-related or otherwise, must rest upon a particular, preexisting law making the person's actions a crime. Typically, this preexisting law is a statute passed by Congress or a state legislature. If no specific, preexisting law is broken, there can be no criminal conviction.

The purpose for this rule is to protect individual citizens from overzealous prosecution. Prosecutors and courts should not be able to make up new criminal laws to punish actions after they have been committed.

On account of this requirement for preexisting law, Congress and state legislatures have in recent years enacted new laws making clear which computer actions constitute crimes.

Let's consider some specific laws that criminalize computer abuse.

The federal Computer Fraud and Abuse Act is a criminal law that punishes people who intentionally cause harm by accessing computers without authority. The legal citation to the act is 18 United States Code Section 1030. The act generally forbids people from knowingly gaining unauthorized access to a computer of the U.S. government or a financial institution or a computer that is used for interstate or foreign commerce (which embraces many computers on the Internet), if that access leads to

- ◆ Classified or national security-related information
- ◆ Records of a financial institution

NOTE

Specific, Preexisting Law Required

The government can't convict a suspect for a crime where there is no specific, preexisting law stating that the suspect's action is a criminal offense. For example, the Philippine government struggled to find a criminal law for prosecuting the author of the "I Love You" virus in 2000 because the country did not have a law specifically criminalizing actions such as the propagation of a computer virus.

NOTE

Unauthorized Access Banners A banner warning that unauthorized access to a network is forbidden can help provide proof that a hacker intentionally committed a crime.

Such a banner might read, for example:

"This is a U.S. government computer system. Government computer systems are provided for the processing of official U.S. government information only. All data contained on government computer systems is owned by the U.S. government and may be monitored, intercepted, recorded, read, copied, or captured in any manner and disclosed in any manner, by authorized personnel. **THERE IS NO RIGHT OF PRIVACY IN THIS SYSTEM.** Systems personnel may give to law enforcement officials any potential evidence of crime found on this U.S. government system. USE OF THIS SYSTEM BY ANY USER, AUTHORIZED OR UNAUTHORIZED, CONSTITUTES EXPRESS CONSENT TO THIS MONITORING, INTERCEPTION, RECORDING, READING, COPYING, or CAPTURING and DISCLOSURE. **IF YOU DO NOT CONSENT, LOG OFF NOW.**"

- ◆ Government records
- ◆ Information on a computer involved in interstate commerce
- ◆ An effect on the government's use of the computer
- ◆ Fraud
- ◆ Damage
- ◆ Trafficking in passwords
- ◆ Extortion

The text of the Computer Fraud and Abuse Act appears at <http://www4.law.cornell.edu/uscode/18/1030.html>. You should study the words of this important law.

Most states also have laws that criminalize unauthorized access to computers.

The federal Wiretap Act, 18 United States Code Section 2511, is a criminal law that punishes unauthorized interception of electronic communications in transit. You can find it at <http://www4.law.cornell.edu/uscode/18/2511.html>.

If the Wiretap Act covers the interception of email while being transmitted, what should a companion law cover to protect email in all stages of its possession by service providers? The federal Electronic Communications Privacy Act, 18 United States Code Section 2701, is a criminal law that forbids unauthorized people from accessing or damaging electronic messages in storage. The text of the law is available at <http://www4.law.cornell.edu/uscode/18/2701.html>.

The key to an action being punishable as criminal is that the suspect *intentionally* do something wrong. Without intent to do something wrong, there can be no crime.

It is easier to show that intrusive hackers acted with wrongful intent if they were notified in advance that they were not authorized to access a system. Therefore, it is wise practice to post banners on network resources warning that access beyond a certain point or in a certain forbidden way is illegal. Such a banner would, for example, warn hackers that they are crossing a legal boundary if they attempt to break into a server.

What is necessary to convict a person suspected of violating a computer crime law? Can a court convict a person and send him to jail on the basis of mere suspicions? Let's consider computer security incidents and the methods by which those incidents can lead to conviction in court.

COMPUTER SECURITY INCIDENTS

Identify and plan for security incidents.

Common types of computer security incidences include viruses, exploratory probes, active intrusions, malicious destruction of data, and denial-of-service attacks. The motives behind these incidents define the major categories of computer attacks, as shown in the following:

- ◆ **Military and intelligence**—Attacks in which spies attempt to learn government secrets or disrupt government operations.
- ◆ **Business**—Attacks in which competitors try to hijack trade secrets.
- ◆ **Financial**—Attacks in which criminals try to trick banks or other financial institutions into sending them money or allocating them credit in an account against which they can make payments.
- ◆ **Terrorist**—Attacks in which politically motivated agents attempt to scare or harm the public by corrupting the computers of government, utilities, or corporations.
- ◆ **Grudge**—Attacks in which disgruntled employees seek revenge on employers by wrecking their information systems.
- ◆ **Consumer fraud**—Attacks in which con artists steal personally identifiable information about consumers (such as Social Security numbers or credit card numbers) so they can impersonate those consumers when purchasing goods or applying for credit or in which the con artists sell consumers bogus goods or services.
- ◆ **Fun**—Attacks in which hackers get thrills, publicity, or payment for breaking into corporate or government computer systems. For some young enthusiasts, the challenge of burglarizing a high-profile system is an intellectual game.

The source of this list is “Fighting Computer Crime” by David Icove, Karl Seger, and William VonStorch (<http://www.cs.nsu.edu/others/seminar/notes/crime1.html>). Refer to it for more information.

Factors that can deter computer crime include prevention measures, such as internal control or access control systems, and detection measures, such as auditing of system activities and supervision of system users.

Damages from computer security breaches can include system downtime, lost employee productivity, wasted effort by system administrators, stolen products or money, physical damage to property, and bad publicity. Often, breaches are crimes. How should an enterprise respond to security breaches? Do some incidents warrant more of a response than others?

The following sections investigate the answers to these questions by looking at how advance planning, computer crime investigation, and legal evidence provide response solutions to incidences.

Advance Planning

A critical element to incident response is to establish an incident plan in advance. Advance planning allows the establishment of priorities, the training of employees before a crisis hits, and the best preservation of legal evidence. Among the steps needed in an advance plan of action are

- ◆ Centralize management of the attack so all of the response can be coordinated.
- ◆ Designate a single person to receive and analyze reports of suspicious or abnormal activities.
- ◆ Make a list of whom to notify.
- ◆ Set procedures for identifying, analyzing, and responding to the attack.
- ◆ Decide how and when to escalate the response to an attack if it grows worse.
- ◆ Designate who has responsibility for which tasks and who within your organization is to be kept informed and mobilized.

- ◆ Specify how to log records of the event and preserve evidence.
- ◆ Establish priorities if there is a tradeoff between preserving evidence and keeping systems in production.
- ◆ Become familiar with the relevant law enforcement authorities and information sharing organizations in advance, and determine which ones to notify at which time.
- ◆ Recognize that a security incident could be more than a technical matter and might warrant coordination with public relations people, corporate attorneys, human resources (if employees are involved), and upper management.
- ◆ Reevaluate security, personnel, and the incident response plan after particular incidents occur. Plans should be regularly reviewed and updated.

Industries and governments have various organizations to collect and disseminate information about computer attacks. Two such organizations are InfraGard (<http://www.infragard.net/>) and Internet Storm Center, at the SANS Institute, (<http://www.incidents.org/>).

Some corporate and government information systems are under attack constantly. How does management decide which attacks to report? Generally, the events to be reported are those that have a substantial impact on an organization (such as damage to assets or reputation) or that are unusual and noteworthy.

See the CIO Cyberthreat Response and Reporting Guidelines posted at <http://www.cio.com/research/security/response.html> for more information on reporting of incidences.

Computer Crime Investigation

One part of the response to a computer security incident can be a computer crime investigation. The investigation might be conducted by private investigators, law enforcement, or a combination of the two. The objective is to minimize risk, while gathering and securing reliable evidence that could be used in a criminal trial. The key to success is the execution of a logical, disciplined plan of action. Step By Step 9.1 outlines the procedures to follow in a computer criminal investigation. Note also that a computer crime investigation can include a more detailed computer forensic investigation, which is discussed in the “Computer Forensics” section later in this chapter.

NOTE

Coordinating with Other Functions

Within the Organization Response to a security incident might require coordination with nontechnical officers within the organization. Public relations staff might need to manage how the incident is reported to the media or customers. If an employee is involved in the incident, the human resources department might need to guide the process by which the employee is confronted and the incident is documented. Upper management might need to be alerted so the necessary resources are made available.

STEP BY STEP

9.1 Classic Computer Crime Investigation from the Perspective of Security Professionals in an Enterprise

1. First, you must detect the intrusion. Detection might come from suspicious or abnormal activity spotted through accidental discovery, audit trail review, or security-monitoring software.
2. Next, you must do whatever is necessary to avoid any additional damage and cut off the potential for liability, such as liability to trading partners who stand to be damaged by the incident.
3. Report the incident to management, being careful to limit knowledge of the investigation and use secure channels of communication.
4. Next is the preliminary investigation, in which you assess damage, witnesses, and whether a crime has occurred and determine what the investigation will need going forward.
5. Next, decide whether disclosure of the incident to government or the media is desired or required. It might be mandatory, for example, to disclose bank fraud to banking regulators.
6. Decide on a course of action, such as tightening of security, maintaining surveillance, or seeking prosecution. The victim enterprise might decide not to pursue prosecution or further investigation because they can be expensive, disruptive, and even embarrassing.
7. Next, assign responsibility for conduct of the investigation, whether it is to internal staff, external consultants, or law enforcement. Issues to consider are cost, investigation control, legal obligations and objectives, and the risk that information about the incident will leak. A possible advantage to using a private investigator rather than law enforcement is that law enforcement often must obtain search warrants (issued by a court) to support its searches and seizures. However, law enforcement can possess greater search and investigation capabilities.

If a search warrant is required, law enforcement must show a court that probable cause exists to believe that a crime has been committed and a search/seizure is needed to investigate.

8. Pinpoint potential suspects (insiders, outsiders, or a conspiracy of both) and potential witnesses, and designate who should interview witnesses.
9. The next step is to plan and prepare for seizure of target systems, including the possible need for special experts and a search warrant. The investigator will want to know as much as possible about the target system in advance to ensure she is properly prepared with team members and equipment.
10. Designate a search and seizure team, including a lead investigator, IT security specialist, legal advisor, and technical staff.
11. Evaluate the risk to the target system before seizing it, including anticipated reaction of the suspect and the risk that evidence will be destroyed.
12. Execute the seizure plan. Secure and search the location, preserve evidence, record each action (such as in a notebook), videotape the process, photograph the system configuration and monitor display, and move the system to a secure location.
13. The final step is to prepare a detailed report documenting facts and conclusions.

The source for the previous list appears at www.cccure.org/Documents/Ben_Rothke/Law-Invest-Ethics.ppt. Refer to it for more detail.

LEGAL EVIDENCE

Explain the laws of evidence and introduce techniques for obtaining and preserving computer evidence.

NOTE

The Cost of Investigating and Prosecuting

As a victim enterprise evaluates the cost of a crime investigation, it should remember that the cost can involve more than what occurs in the investigation itself. After the investigation gathers evidence, it can lead to criminal proceedings in court. The proceedings are normally a discovery phase, a grand jury phase, and a trial phase. The court proceedings can require the production of documents and collection of further evidence, as well as the delivery of testimony. All these can consume considerable employee time and might, as a practical matter, be a deterrent to reporting in the first place.

One possible objective of incidence response is to gather evidence for prosecuting a criminal perpetrator. Evidence might be used by criminal investigators to find a perpetrator. Investigators might also use it to convict the perpetrator of violating a criminal law, such as the Computer Fraud and Abuse Act introduced in the “Criminal Law and Computer Crime” section.

Evidence is also used in court to resolve civil litigation such as contract, property, or tort disputes.

Evidence is anything that demonstrates a point to a court or persuades the court that a fact is true. Evidence can include a document or a log of activity on a network (*documentary* evidence); testimony from a witness about his direct observation of something; tangible objects (*real* evidence); or models, illustrations, or simulations (*demonstrative* evidence).

The rules of evidence govern which evidence can be admitted into court to demonstrate something, such as the fact that a criminal defendant logged onto a server at a certain time of day. The basic rule is that any “relevant” evidence can be admitted and considered by the court. Evidence is relevant if it tends to answer a question before the court (for example, whether the defendant logged on at that time).

Credibility or Weight of Evidence

Some evidence is stronger or more credible than other evidence. The credibility of evidence is usually determined by the trier of fact—in other words, the judge or jury in the court.

Strong evidence of a fact is called *direct* evidence; weaker evidence is called *circumstantial* evidence. Circumstantial evidence requires the trier of fact to leap through more logical inferences to conclude that the fact supported by the evidence is true. For example, a recipient’s complete, unmodified log of received emails is direct evidence of what email the recipient received. But it is only circumstantial evidence of what a particular sender might have sent. To determine from the log what was sent requires knowledge of the system between sender and receiver and the drawing of inferences from that.

Both direct evidence and circumstantial evidence can be admitted as evidence in court (provided they otherwise satisfy the rules of evidence), but direct evidence carries more weight. In other words, direct evidence is more believable.

To help keep weak evidence out of court, the rules of evidence hold that evidence must be established as authentic, not hearsay, and compliant with the best evidence rule. The concepts of authenticity, hearsay, and best evidence rule should be understood more as rules of thumb rather than hard rules that are followed slavishly in court.

Regardless of these technical rules, there is a practical aspect to evidence. Evidence that is offered or supported by credible witnesses and professional investigative techniques is much more likely to win the day in court. A systematic, disciplined method for gathering evidence is persuasive to the trier of fact.

Proof of Authenticity

To be *authentic*, evidence must be supported by something showing that the evidence is what it purports to be. Proof of authenticity need not necessarily be extremely strong to support admission in court. In other words, for admissibility purposes, proof of authenticity does not necessarily require military-grade security. But if proof of authenticity is weak, the trier of fact might assign the evidence little or no weight.

Hearsay

The “hearsay rule” excludes from court a statement made outside the court that is repeated for the purpose of showing the statement is true. For example, a letter from Jane that says, “Bill bought a car in July,” is hearsay if it is offered in court as evidence that Bill did buy a car in July. However, the hearsay rule has many exceptions. One of those exceptions is that records kept in the ordinary course of business are admissible even though they are hearsay. Very often, business computer records are admitted into court (even though they are technically hearsay) because they were created in the ordinary course of business. Creation of records in the ordinary course of business implies a disciplined, logical method to record-making.

Best Evidence Rule

The “best evidence rule” says that to prove the terms of a “writing,” the original writing must be produced in court—not a copy—because the original is more reliable. But the best evidence rule has many exceptions, and in the electronic realm the rule is confusing.

NOTE**Professionalism in Gathering**

Evidence The evidence that is most powerful in court is that which is captured in a logical, controlled fashion.

NOTE

Segregation of Duties Makes for a Good Chain of Evidence The famous case *United States v. Poindexter* (Crim. No. 88-0080-1) (D.D.C. 1990) illustrates the use of computer evidence in court. The evidence consisted of records of email in a closed, local area network. The records were stored on magnetic tape, under the supervision and custody of the network administrator. The court admitted and relied on the records, but only after the administrator testified about the reliability of the system and the controls in place to protect the records. The administrator was a neutral party and therefore had duties that were *segregated* from the people who created and relied on the email in question. He established that the tapes stayed under his control (locked in his office) and therefore that a good chain of evidence supported the records.

NOTE

Can Imperfect Evidence Be Used in Court? Even though the systematic, disciplined gathering of evidence is most persuasive to courts, *imperfect evidence* can still be helpful to investigators—even if it proves to be inadmissible in court.

Further, good trial lawyers can sometimes find creative and surprising ways to use evidence in court, even though it might be subject to criticism because it could have been fabricated or came from an imperfect chain of evidence. For example, a lawyer might offer as evidence in court a log showing that someone used Mary's password to access a system. The lawyer might offer this evidence to prove that Mary herself accessed the system, even though it is possible a hacker had stolen Mary's password.

When an electronic writing is at issue, you can most easily satisfy the best evidence rule with respect to that writing by persuading the court that the evidence being offered is an accurate representation of the writing.

The best evidence rule should not be understood as requiring that the best or most direct evidence be admitted in court. However, as stated previously, direct evidence does carry more weight than circumstantial evidence.

Chain of Evidence

Controls are practical measures that reduce the chance records are changed or corrupted. Examples of controls are *audit trails* and *segregation of duties*. Audit trails are detailed records of a process, showing what happened, when, where, and how. Segregation of duties means having one person in charge of one part of a record-making process and having an independent person responsible for another part of the process. The presence of better controls makes computer records more believable. Controls denote logic, discipline, and accuracy.

One form of control is a *chain of evidence* (also known as *chain of custody*). The chain of evidence is a series of records showing where evidence came from, who was responsible for it, what happened to it, how it was protected, whether it was changed, and so on. A good chain of evidence also includes procedures to ensure evidence is not lost or corrupted. When an investigator creates a chain of evidence, his objective is to ensure he can account for possession and integrity of the evidence from its origin to the time it is brought into the courtroom.

There is no perfect way to obtain and preserve evidence, and there is no perfect form of evidence. In a thorough investigation, the more evidence the better, even if some of it is imperfectly collected and preserved.

Records created according to routine business procedures, under strong internal controls, and then protected through a good chain of evidence are of higher credibility and value.

The Fourth Amendment

The Fourth Amendment to the U.S. Constitution protects citizens from unreasonable searches and seizures by government. Therefore, law enforcement normally needs a court-issued warrant before searching or seizing evidence, although there are exceptions, such as when evidence is in plain view.

Issuance of a warrant usually requires showing a judge that law enforcement has probable cause to believe the evidence is relevant to a crime. After a warrant is issued, the search for evidence should stay within the terms of the warrant. If law enforcement believes more evidence is available, it should obtain a warrant for that additional evidence. Under the exclusionary rule, evidence obtained in violation of the Fourth Amendment is excluded from court. The purpose of the exclusionary rule is to penalize law enforcement if it violates the Fourth Amendment.

COMPUTER FORENSICS

Introduce techniques for obtaining and preserving computer evidence.

Forensics is the use of science and technology to investigate and establish facts that can be used in court. When using forensics for computer incidents, the one objective is to preserve evidence from the earliest moment possible.

Collection and preservation of evidence is best performed by forensics experts with special training. Consider calling in outside experts.

Still, staff who are not forensics experts can aid an investigation by keeping a disciplined, detailed journal of what happened during an incident and when the events occurred. Secure files that log activities on a network can be powerful evidence for use in investigations and court. The more extensive the logs, the better because extensive logs signify discipline and diligent effort. Ideally, the logs would be maintained all the time, not just in response to an incident. (By maintaining them all the time, you increase the chance a court will view them as routine business records that are exempted from the hearsay rule, which was discussed earlier in the chapter in the “Legal Evidence” section.) The logs are more credible if their integrity is protected with such measures as digital signatures; secure time stamps; segregation of duties; and the use of dedicated, separate computers.

NOTE

Computer Forensics to Assess Email Evidence

Suni Munshani v. Signal Lake Venture Fund II, LP

(Massachusetts Superior Court, Civil Action No. 00-5529 BLS) demonstrates the use of computer forensics in a dispute over Internet email. Plaintiff Munshani sued the defendant company claiming that the company's CEO promised to grant him warrants for purchase of stock at a favorable price in exchange for the plaintiff's work for the company. To support his claim, the plaintiff produced an email record purporting to make the promise. The defendant, on the other hand, proved the email record was fake by producing a thorough forensic analysis of the plaintiff's and defendant's email logs. The analysis showed that the plaintiff's record was an alteration of an authentic email. Anomalies in the email headers, together with a date stamp that was five months too late, showed the plaintiff's record to be a forgery. See the court order, the forensic report, and explanatory articles at <http://www.signallake.com/litigation>.

The source of this information is the article “Email Tampering, This Time the Good Guys Won,” by M. Weingarten and A Weingarten, which appeared in the January 2002 issue of *Business Communications Review*.

NOTE

Practical Forensics For more information on the practical use of computer forensics, see Illena Armstrong's article "Computer Forensics, Tracking Down the Clues," which appeared in the April 2001 issue of *SC Magazine* (http://www.scmagazine.com/scmagazine/2001_04/cover/cover.html).

NOTE

Best Practices For more information about how to seize computer evidence, see the article "Best Practices for Seizing Electronic Evidence: A Joint Project of the International Association of Chiefs of Police and the U.S. Secret Service" at www.treas.gov/usss/electronic_evidence.htm.

When collecting evidence about a particular incident, a single individual should be designated to coordinate the entire process and ensure that all procedures are followed. A detailed, chronological notebook should be kept of all steps followed to collect and transport evidence. Tamper-proof copies of evidence should be made by properly trained personnel, using competent tools. Evidence should be sealed, tagged, and logged into the incident notebook. Evidence must be stored in a secure location, and every time the evidence is moved or examined, details should be recorded in the evidence notebook. These efforts are the earmarks of a disciplined, credible effort to gather evidence.

Even when a company calls law enforcement to collect evidence, the company should have its own private investigators making copies of evidence in case it is needed for private litigation or insurance claims.

The techniques for seizing and preserving electronic evidence so as not to alter or destroy it follow:

- ◆ Restrict physical and remote access to the computer.
- ◆ If computer is off, do not turn it on.
- ◆ If computer is on, photograph the image showing on the screen and then unplug the computer.
- ◆ Do not touch the keyboard.
- ◆ Do all forensic analysis of the electronic evidence from a mirror copy of the disk on which the evidence is originally stored.
- ◆ Don't trust the subject computer's operating system; conduct analysis on a copy using the operating system of a trusted computer.

Step By Step 9.2 outlines the techniques you should use to examine a PC.


STEP BY STEP

9.2 PC Examination Checklist

1. Before starting a computer forensics examination, get appropriate authority from corporate management. If the investigator is in law enforcement, a court-issued search warrant might be necessary.
-

2. If the machine is on, turn it off by pulling the plug. To record the state of the computer before it was unplugged, photograph the image displayed on the monitor.

3. Before moving the computer, document the hardware configuration with photographs and tags on cables, as shown in Figure 9.1. Collect, package, and label removable media such as floppy disks, tapes, and CDs present in the premises of the PC.

**FIGURE 9.1**

A careful forensic investigator photographs the system's location and general setup before moving the computer.

continues

NOTE

Best Practices For more information about how to handle the examination of a PC, see the article “Best Practices for Seizing Electronic Evidence: A Joint Project of the International Association of Chiefs of Police and the U.S. Secret Service” at www.treas.gov/usss/electronic_evidence.htm.

continued

4. Transport the computer to a secure location.
5. Boot the computer without booting from the suspect hard drive itself. Boot from a floppy, or remove the hard drive and examine it using a separate computer dedicated to forensic examination.
6. Using forensic software, make a bit-stream image of the suspect drive; then run a hash of the suspect hard drive and the image to confirm the data in the two are the same. Next, document the system date and time. Forensics software can then be used on the image copy to run keyword searches through files, free space, and slack space. Popular forensic software packages include AccessData Development’s Forensic Toolkit (FTK), Guidance’s EnCase, and NTI’s SafeBack.

It is better to analyze a mirror image of the contents on a drive than the contents actually on the drive. By analyzing the mirror image, the forensic investigator avoids altering the original data.

For more information on the elements in Step By Step 9.2, see the following:

- ◆ “Digital Forensics: Crime Seen,” an article by Bill Betts that appeared in the March 2000 issue of Information Security Magazine (<http://www.infosecuritymag.com/articles/march00/cover.shtml>).
- ◆ “Legal Aspects of Collecting and Preserving Computer Forensic Evidence,” an article by Franklin Witter that appears on the Web site <http://rr.sans.org/incident/evidence.php>.

Step By Step 9.3 shows you the steps a computer forensic expert should take when analyzing what is on a computer.

STEP BY STEP

9.3 The Steps of a Computer Forensic Analysis

1. Make a bit-level image copy of the suspect disk.
 2. Make a cryptographic hash or digest of the disk as a whole and all directories, files, and disk sectors.
 3. Perform analysis in a secure environment.
 4. Use forensics software to find hidden, deleted, or encrypted files.
 5. Boot the suspect system with a trusted operating system. Run a complete system analysis.
 6. To discover any background or malicious programs and learn of any system interrupts, reboot the suspect system with its original operating system.
 7. Examine backup media, such as CDs or floppies.
 8. Investigate any files that are protected with passwords or encryption. Techniques such as password crackers and interviews of suspects can lead to the opening of files.
-

The list in Step By Step 9.3 is drawn from www.cccure.org/Documents/Ben_Rothke/Law-Invest-Ethics.ppt, and more information can be found there.

COMPUTER ETHICS

Discuss computer ethics.

What is the relationship between criminal law and ethics? Should the principles stating what is and is not criminal be similar to the principles of what is and is not ethical? Recall the Computer Fraud and Abuse Act discussed earlier. Compare it to the Request for Comments (RFC) 1087 titled “Ethics and the Internet,” published January 1989 by the Network Working Group of the Internet Activities Board.

518 Part I EXAM PREPARATION

RFC 1087 declares unethical and unacceptable any activity which purposely

- ◆ Seeks to gain unauthorized access to the resources of the Internet
- ◆ Disrupts the intended use of the Internet
- ◆ Wastes resources (people, capacity, computer) through such actions
- ◆ Destroys the integrity of computer-based information
- ◆ Compromises the privacy of users

How similar are these principles to those in the Computer Fraud and Abuse Act? Notice that both the principles and the act warn against unauthorized access to computers that leads to some kind of injury.

To whom does RFC 1087 apply? It applies to all Internet users, which includes computer security professionals, but it also includes many more people.

The computer security profession aspires to have its members recognized as trustworthy and credible. How might that aspiration be achieved? (ISC)² publishes a code of ethics that is specific to computer security professionals and maintenance of their professional certification. The (ISC)² Code of Ethics, which is published at <http://www.isc2.org/cgi/content.cgi?category=12>, requires CISSPs to

- ◆ Protect society and infrastructure
- ◆ Act honestly and legally
- ◆ Deliver competent professional service
- ◆ Uphold the profession

Breach of this code can lead to revocation of CISSP certification.

The (ISC)² Code is written in the form of four general cannons (which are mandatory), followed by explanatory guidance.

Notice that because the (ISC)² Code requires a CISSP to uphold high ethical standards, the CISSP would normally be expected to abide by RFC 1087 when the professional is on the Internet.

One of the guidelines in the (ISC)² Code requires that a CISSP avoid conflicts of interest. A conflict of interest occurs when a professional owes loyalty to two different people who have competing interests, such as the professional's employer versus a vendor to the employer or the employer versus the professional's own self interest. For example, a computer security professional has a conflict of interest if her employer asks her to investigate the presence of gambling over the employer's information systems when the professional is one of those who has in fact been participating in the gambling activities.

Computer ethics should be promoted within organizations through training and published reminders to end users. Employee manuals should include material on computer ethics.

NOTE

Study Ethics Code You should study the (ISC)² Code of Ethics thoroughly. It is not written as a black-and-white set of detailed rules, but rather as general principles intended to promote good ends, such as professionalism, truthfulness, and safe computing practices.

CASE STUDY: CROSS-EXAMINING THE FORENSICS EXPERT

ESSENCE OF THE CASE

Following is a list of key points that make up the essence of this case:

- ▶ A computer forensics expert examined a suspect's computer.
- ▶ The suspect's attorney is probing for shortcomings in the expert's work that would suggest it is not worthy of credibility.

SCENARIO

Shannon testifies in court about the computer forensics techniques she used as a private investigator. Shannon's client, Consolidated Engineering, feared that one of its former engineers, David Smith, had stolen secret drawings for a new product. Lawyers for Consolidated had succeeded in obtaining a subpoena requiring Smith to allow his home computer to be inspected.

Shannon led the inspection and claims she discovered on Smith's hard disk copies of drawings belonging to Consolidated, with time stamps showing they were written to the disk after Smith's departure from Consolidated.

You are the attorney for Smith. You want to discredit Shannon's testimony. What questions would you ask her on cross-examination?

continues

CASE STUDY: CROSS-EXAMINING THE FORENSICS EXPERT*continued***ANALYSIS**

The attorney should probe whether Shannon is a competent professional and a trustworthy witness. It might be that Shannon planted the drawings on Smith's machine to frame him. These are the types of questions the attorney might ask:

- ▶ Did Shannon have incentive to fabricate the evidence? Does she have a reputation for being ethical and credible?
- ▶ Did Shannon have a separate witness to work with her and observe and document her actions as she inspected Smith's computer?
- ▶ Did Shannon preserve the evidence with a chain of evidence showing who controlled and protected the evidence at all times starting from when she first touched the machine in question?
- ▶ What techniques did Shannon employ to prevent alteration of the data during and after inspection? Did Shannon work from a mirror image of data from Smith's hard disk, or did she work from the original disk directly?
- ▶ How did Shannon ascertain whether the system clock on Smith's computer was set to properly time stamp files?

CASE STUDY: PROVING COPYRIGHT INFRINGEMENT**ESSENCE OF THE CASE**

- ▶ Bill's employer suspects a thief is stealing its proprietary data.
- ▶ The thief is encrypting its data.
- ▶ Is it legal and ethical for Bill to intercept the thief's data and break the thief's encryption?

SCENARIO

Bill is a CISSP employed by XYZ Music, an online broadcaster of live concerts. XYZ suspects that Loco Music has found a way to break the encryption XYZ uses to scramble its broadcasts and capture the content so that Loco can resell it as an encrypted product to a small group of elite clients. But XYZ has no proof that Loco is doing this.

Bill knows how to break Loco's encryption. He suspects that if he taps into Loco's Internet transmission and breaks its encryption, he will have proof that Loco is stealing content from XYZ. Bill plans to log the results as evidence. Is Bill about to embark on a wise plan of action?

CASE STUDY: PROVING COPYRIGHT INFRINGEMENT

ANALYSIS

Bill is about to venture into dangerous waters. Although Loco might be infringing XYZ's copyright and might be violating the Digital Millennium Copyright Act, Bill does not know that. What's more, Bill himself will be at risk of infringing Loco's copyright and of violating the DMCA. When he breaks Loco's encryption, he might be defeating a security measure that Loco applies to protect its own copyrighted material, some or all of which might legitimately be owned by Loco.

Bill should be careful about "tapping" into Loco's transmission. If, for example, he goes to a server owned by Loco and accesses the transmission without authority, he might be violating the Computer Fraud and Abuse Act, the Wiretap Act and state computer crime laws, as well as RFC 1087's ethical teaching that Internet users are not to seek unauthorized access to Internet resources.

As a CISSP, Bill has an ethical duty to avoid unlawful professional conduct.

CHAPTER SUMMARY

It's hard to predict precisely what legal and investigation material will be covered on the exam. Technology, law, and methods are changing, and even experts can disagree on what is right, what is wrong, what is important, and what is not important. It is hoped that you gain an intuitive sense of the subject by studying this chapter and the materials cited in it.

This chapter introduced the intellectual property concepts of patent, copyright, and trade secret and explained that serious copyright and trade secret violations can be crimes. It identified other key American computer crime laws: the Computer Fraud and Abuse Act, the Wiretap Act, the Electronic Communication Protection Act, and the Digital Millennium Copyright Act.

The motivations for and responses to computer attacks were introduced. The key to good response to an incident is to have a plan in place in advance, so procedures, contacts, and priorities don't have to be worked out in a crisis.

A prime objective of a computer crime investigation is to collect and preserve legally useful evidence. Organization, logic, and thorough documentation are the qualities that will win the results of an investigation favor in court.

KEY TERMS

- Authenticity
- Best evidence rule
- Chain of evidence or chain of custody
- Conflict of interest
- Copyright
- Digital Millennium Copyright Act
- Directive on data protection
- Exclusionary rule
- Fair information practices
- Forensics
- Hearsay
- HIPAA

continues

CHAPTER SUMMARY *continued*

- Gramm-Leach-Bliley
- License
- Mirror image
- Patent
- Privacy
- Safe harbor on data protection
- Trade secret
- U.S. Fourth Amendment

Although they can be applied in flexible and surprising ways, the rules of evidence structure and limit the use of evidence in court. Evidence gathered in a disciplined, methodical way is more credible. A critical technique for adding to the value of computer evidence is a good chain of evidence, which documents where evidence comes from, whether it was changed, and who had custody of it.

When collecting computer evidence, law enforcement should be careful to get proper search warrants, lest it violate the U.S. Fourth Amendment guarantee that citizens will be free from unreasonable searches and seizures. Violation of the Fourth Amendment can lead to the exclusion of evidence from court.

Good computer forensics techniques discover hidden evidence and avoid altering or destroying any evidence.

Computer security professionals are expected to uphold high ethical standards. This makes them more credible as witnesses in court and more trustworthy as stewards of information resources.

APPLY YOUR KNOWLEDGE**Exercises****9.1 Connecting the Key Principles**

Reread this chapter, and look for the key philosophical principles that apply to each of the topics covered here. Notice the interrelationships between the principles in each of the topics. Write sentences describing the interrelationships you see; the process of writing will help you remember as you prepare for the exam.

Estimated Time: 30 minutes

Answer to Exercise 9.1:

1. Notice how computer crime law is based on ethical principles of good computer practices.
2. Also note how the purpose of evidence law is to find credible representations of fact, and the evidence of computer activities that is most credible is that which is gathered according to disciplined, methodical procedures.
3. The best forensic techniques emphasize logical, controlled steps for securing evidence and memorizing it in records.
4. Notice that privacy is achieved by following logical, disciplined steps to notify individuals about how their private information will be used. Privacy is about being honest and truthful, which are ethical qualities expected of CISSPs.
5. Finally, you should have learned how third parties can promote desired results in information management. Segregation of duties makes records more credible. And privacy is protected by requiring law enforcement to seek approval from an independent third party (that is, a court) before a search of private information is conducted.

Review Questions

1. What factors should be considered before a computer security incident occurs?
2. What are some leading laws requiring businesses to secure their information resources?
3. How does a company protect its rights to trade secrets?
4. What are the prerequisites to prosecuting a suspect for a crime?
5. What are the essential provisions of the Computer Fraud and Abuse Act?
6. What are the key ethical principles for a computer security professional?
7. Identify basic principles of fair information practice.
8. How does one make a chain of evidence?

Exam Questions

1. Which of the following is *not* always required for the government to secure a criminal conviction of a suspect?
 - A. A confession signed by the suspect
 - B. Evidence that the suspect broke a criminal law
 - C. A specific law stating that the act committed by the suspect was a crime
 - D. Evidence that the suspect acted with intent

APPLY YOUR KNOWLEDGE

2. A police officer suspects Joe is using his computer to break into Acme, Inc.'s corporate information systems. The officer seizes Joe's computer and conducts a careful forensic analysis of the data stored on Joe's hard drive. Later, when Joe is being prosecuted in court, the judge determines that the police officer should have obtained a search warrant before seizing and searching Joe's computer. What is the judge likely to do?
 - A. Convict Joe of violating the Computer Fraud and Abuse Act.
 - B. Conduct his own forensic analysis of Joe's computer.
 - C. Exclude from court the evidence obtained by the police officer from Joe's computer.
 - D. Levy a fine against Acme, Inc.
3. Armed with a warrant for searching and seizing a suspect's computer, a police investigator enters a suspect's home and prepares to seize his computer for further investigation. The computer is turned on. What should the investigator avoid doing?
 - A. Photographing the computer
 - B. Tagging the cables coming from the computer so the investigator can remember which cable was plugged into which port
 - C. Shutting down the computer's operating system
 - D. Removing the computer to the investigator's facilities for careful analysis
4. Which is least likely to be an ethical violation?
 - A. Under the direction of the CEO, a security manager destroys records of the CEO's wrongdoing.
 - B. A security manager says she will advocate that her company purchase a certain security product if the vendor sponsors her vacation on a cruise ship.
 - C. A security manager, in accordance with his company's published policy, reviews the content of employee email on company servers.
 - D. A security manager misleads a journalist to protect her company's interests.
5. Which of the following is least likely to be a crime?
 - A. Imitating a new competitor's business strategy
 - B. Selling pirated music
 - C. Stealing a competitor's secret method for organizing a database
 - D. Exceeding authority on public ISP servers to view private email records
6. An IS employee on duty Sunday night discovers an unfolding computer security incident. What would be the best source of information on what the employee should do?
 - A. A leading textbook on computer security
 - B. The Computer Fraud and Abuse Act
 - C. The FBI
 - D. An incident response plan previously established by the employee's management
7. Which is typically *not* part of a computer forensic investigation?
 - A. Making a mirror image of a subject computer's hard disk
 - B. Erasing corrupted files

APPLY YOUR KNOWLEDGE

- C. Searching for hidden data in slack space or attached to the end of files
 - D. Moving a subject computer to the investigator's office
8. After a security incident begins, you set up a facility for logging data as evidence of what is happening. After you start the logging process, you think of a way in which a clever hacker could defeat or corrupt the logged data. Which is the better course of action?
- A. Preserve the log as is.
 - B. Destroy the log.
 - C. Obtain advice by submitting an inquiry to the (ISC)² ethics committee.
 - D. Notify the Internet Storm Center at the SANS Institute (<http://www.incidents.org/>) of how the log might be corrupted.
9. Which of the following is *not* part of a typical chain of computer evidence?
- A. Making a mirror image of data on a hard disk
 - B. Storing data media in protective bags, labeled with date, time, place of origin, and identity of custodian
 - C. Videotaping the installation of a new PC
 - D. Detailing in a notebook the methods used to collect, protect, and store data

Answers to Review Questions

1. Before a security incident occurs, advance planning and training are critical. The plan should address your organization's priorities and the tradeoffs between the collection of evidence for prosecution and the maintenance of systems in production. The plan should address whom to notify and when. For more information, see the section "Advance Planning."
2. The following are laws requiring information security on the part of corporations: the Foreign Corrupt Practices Act, the Gramm-Leach-Bliley Financial Modernization Act, and the Healthcare Insurance Portability and Assurance Act (HIPAA). For more information, see the section "Government Regulations."
3. A company that wants to maintain the value of its trade secrets endeavors to keep the secrets a secret. It enters nondisclosure agreements with employees and trading partners who need to know the secrets. It also protects the secrets with encryption and copy controls. For more information, see the section "Trade Secrets."
4. To convict a suspect of a crime, the suspect must have intentionally committed an act that was previously defined by law (normally a statute passed by Congress or a state legislature) as a crime. A prosecutor must produce evidence to a court showing, beyond a reasonable doubt, that the suspect committed the act. For more information, see the section "Criminal Law and Computer Crime."

APPLY YOUR KNOWLEDGE

5. The Computer Fraud and Abuse Act forbids knowing, unauthorized access to a computer of the U.S. government or a financial institution or which is used for interstate or foreign commerce, if that access leads to any of the following: classified or national security-related information, records of a financial institution, government records, information on a computer involved in interstate commerce, an effect on the government's use of the computer, fraud, damage, trafficking in passwords, or extortion.

For more information, see the section "Criminal Law and Computer Crime."

6. These summarize the CISSP's ethical duties: Do protect society and infrastructure; do behave honestly and legally; do deliver professional service; and do uphold the profession. For more information, see the section "Computer Ethics."

7. An individual who is the subject of collection of personally identifiable information should have right to the following: notice about which data will be collected and how it will be used; choice about whether data will be collected; access to collected data; reasonable protections for accuracy, integrity, and security of collected data; and rights to seek redress for abuse of data.

For more information, see the section "Privacy Law."

8. There is no single way to make a good chain of evidence. A chain of evidence is persuasive documentation and procedures that show a court where evidence came from, how it was stored and protected, who stored and protected it, and that it was not tampered with. The chain can include chronological notes in a notebook, secure storage facilities, labels on storage media, time stamps, and employee training. For more information, see the section "Chain of Evidence."

Answers to Exam Questions

1. **A.** To secure a conviction, the government needs proof that the suspect intentionally broke a specific criminal law. A confession can be the proof required. But if the suspect does not confess, the government can prove its case by other means. For more information, see the section "Computer Law and Computer Crime"
2. **C.** When the judge determines that the police officer should have obtained a search warrant in advance, the judge is in effect saying that the officer violated Joe's right under the Fourth Amendment to be free of unreasonable searches and seizures by the government. A typical remedy when the Fourth Amendment has been violated is to exclude from trial any evidence the government obtained through the illegal search and seizure. For more information, see the section "The Fourth Amendment."
3. **C.** When a forensics investigator seizes a computer that he finds turned on, normally the best way to shut down the computer is to unplug it from its power source. Shutting down the operating system can alter or destroy evidence on the computer. For more information, see the section "Computer Forensics."
4. **C.** The manager does not violate the privacy rights of employees by examining their email where the company has told employees (such as through a published policy) that their email is not private. Ethical rules do forbid security professionals from destroying important data (which is dishonest), maintaining a conflict of interest, or lying. For more information, see the section "Computer Ethics."

APPLY YOUR KNOWLEDGE

5. **A.** A company usually has no right to exclude others from copying the way it conducts business. But selling pirated music appears to violate copyright laws. Stealing a secret method appears to be theft of the competitor's trade secret, and viewing email without authority appears to be a violation of the Electronic Communication Privacy Act. For more information, see the section "Intellectual Property Law."
6. **D.** A previously established plan should give the employee the specific instructions she needs for her particular facility and should set the priorities that are important for her enterprise. For more information, see the section "Advance Planning."
7. **B.** A key objective of a computer forensics investigation is to avoid altering or destroying data. For more information, see the section "Computer Forensics."
8. **A.** No evidence is perfect. Better to preserve what evidence is collected than to destroy it. For more information, see the section "Legal Evidence."
9. **C.** Typically, a chain of computer evidence is a series of techniques and procedures for gathering and preserving evidence from a computer that has previously been in use. For more information, see the section "Chain of Evidence."

APPLY YOUR KNOWLEDGE**Suggested Readings and Resources**

1. Hutt, Arthur E., Seymour Bosworth, and Douglas B. Hoyt. *Computer Security Handbook, Third Edition*. John Wiley & Sons, 1995.
2. Mcmillian, Jim, "Importance of a Standard Methodology in Computer Forensics," May 2, 2000. This article is available only on the Web, at this URL: <http://rr.sans.org/incident/methodology.php>.
3. Staggs, Jimmy. "Computer Security and the Law." published by SANS Institute on December 1, 2000. (A copy of the article is available at <http://rr.sans.org/legal/law.php>).
4. Tipton, Harold F., and Micki Krause, eds. *Information Security Management Handbook, Fourth Edition, Volume I*. CRC Press, 1999.
5. Tipton, Harold F., and Micki Krause, eds. *Information Security Management Handbook, Fourth Edition, Volume II*. CRC Press, 2000.
6. Tipton, Harold F., and Micki Krause, eds. *Information Security Management Handbook, Fourth Edition, Volume III*. CRC Press, 2001.
7. Welch, Thomas. "Computer Crime Investigations & Computer Forensics," *Information Systems Security*, Summer 97, Vol. 6 Issue 2, p56. (A copy of the article is also available on the Web at this URL: http://telecom.canisius.edu/cf/computer_crime_investigation.htm).
8. Winn, Jane K., and Benjamin Wright. *The Law of Electronic Commerce, Fourth Edition*. Aspen Law & Business, 2001.

OBJECTIVES

Understand the idea of classifying assets and identifying threats and countermeasures that apply to classes.

- ▶ One of the problems with security assessments is becoming overwhelmed by too much detail. One way to help cope is to deal with classes of things rather than individual assets.

Understand some of the most common vulnerabilities and how they affect different asset classes differently. These include

- **Understand general principles that apply to the theft of information and assets.**
 - **Know the general criteria that apply to the location and construction of facilities.**
 - **Understand basic methods of controlling physical access to an area.**
 - **Know the basic issues relating to regulating the power supply for computers and other equipment.**
 - **Understand common sources of exposure to water and simple countermeasures.**
- ▶ Examining classes of assets and classes of vulnerabilities helps to impose a framework on risk assessment.

Understand some of the most common vulnerabilities and how they affect different asset classes differently.

- ▶ When common vulnerability topics are defined, the threat to specific assets can more readily be addressed. Each threat can be explored, and countermeasures developed to mitigate the threat.



CHAPTER 10

Physical Security

OBJECTIVES

Understand issues and controls related to removable electronic media.

- ▶ Removable media, such as disks and tape, complicates the physical security picture. Not only do computers have to be secured, but we must somehow prevent data from being stolen by preventing removal of the media it resides on.

Understand issues relating to storage of paper.

- ▶ Data that resides on electronic media is not the only type of data at risk. Often, more critical copies of the data lay in printed reports which may be transported out of secured areas, or disposed of without thought for their sensitive nature. In addition, the paper itself may be in need of protection. Checks and other forms which when printed represent monetary value, must be treated differently than other raw paper stocks.

Know the most common issues relating to disposal or erasure of data.

- ▶ Many issues arise with disposal. The most important and rather obvious—and probably most neglected—is that sensitive waste can retain its sensitivity. Simple erasure of computer files might not actually delete data; even if the data is deleted or overwritten, retrieving the data might still be possible with special techniques.

Describe physical intrusion detection methodologies and products.

- ▶ While we all are familiar with alarms, cameras, and guards as solid products, which can alert us to the presence of intruders, and of fences and other inhibiting protection devices, their proper selection and use should be studied.

OUTLINE

Introduction	532
Classifying Assets to Simplify Physical Security Discussions	533
Vulnerabilities	535
Selecting, Designing, Constructing, and Maintaining a Secure Site	538
Site Location and Construction	539
Physical Access Controls	540
Active Physical Access Controls	541
Passive Controls	542
Power	544
Power Issues: Spikes, Surges, and Brownouts	545
Minimizing Power Problems	545
Environmental Controls: Air Conditioning, Humidity, and Temperature	547
Water Exposure Problems	548
Fire Prevention and Protection	549

OUTLINE

Tape and Media Library Retention Policies	553
Document (Hard-Copy) Libraries	555
Waste Disposal	556
Physical Intrusion Detection	559
Chapter Summary	563
Apply Your Knowledge	565

STUDY STRATEGIES

- ▶ Remember that the Common Body of Knowledge is intended to be “abstract and stable” and “independent of necessary skills, tasks, activities or technologies.” When studying, concentrate on general issues (for example, what costs and constraints a card access system imposes as part of a perimeter control strategy) and how to apply specific knowledge, rather than on specifics (for example, characteristics of various types of smart cards).
- ▶ Concentrate on how security issues and measures relate to one another and affect one another. For example, access control card systems affect power supply issues, fire protection, privacy, staffing, and costs as well as the obvious issue of keeping the wrong people out and letting the right people in.
- ▶ Remember that the physical security material in this chapter is part of a broader picture, and concentrate on how these topics relate to material from the other domains.

“The Physical Security domain addresses the threats, vulnerabilities, and countermeasures that can be utilized to physically protect an enterprise’s resources and sensitive information. These resources include people, the facility in which they work, and the data, equipment, support systems, media, and supplies they utilize. The candidate will be expected to know the elements involved in choosing a secure site, its design and configuration, and the methods for securing the facility against unauthorized access, theft of equipment and information, and the environmental and safety measures needed to protect people, the facility, and its resources.”

—Common Body of Knowledge study guide

This chapter covers Domain 10, Physical Security, 1 of 10 domains of the Common Body of Knowledge (CBK) covered in the Certified Information Systems Security Professional Examination. This domain has been divided into several objectives for study.

INTRODUCTION

Physical security refers to the provision of a safe environment for information processing activities and to the use of the environment to control the behavior of personnel.

The objectives in this chapter are explained and supported by observing the categories defined by (ISC)² and addressing a number of supporting topics. (ISC)² groups physical security issues into five categories. These are

- ◆ **Facility requirements**—Such as site selection and construction and perimeter control
- ◆ **Technical controls**—Such as card or token systems
- ◆ **Environmental/life and safety**—Such as power and fire issues
- ◆ **Physical security threats**—Such as weather and other natural events and intentional attacks
- ◆ **Elements of physical security**—Such as sensors and surveillance

This chapter examines the following topics in the realm of physical security:

- ◆ Classifying assets
- ◆ Theft
- ◆ Site location and construction
- ◆ Physical access
- ◆ Power
- ◆ Air conditioning
- ◆ Water exposure and problems
- ◆ Fire prevention and protection
- ◆ Tape and media library retention policies
- ◆ Document (hard-copy) libraries
- ◆ Waste disposal
- ◆ Offsite storage
- ◆ Physical intrusion detection
- ◆ Physical attack parameters

CLASSIFYING ASSETS TO SIMPLIFY PHYSICAL SECURITY DISCUSSIONS

Understand the idea of classifying assets and identifying threats and countermeasures that apply to these classes.

The principles of physical security are no different from those of information security: Identify the assets you need to protect, assess vulnerabilities and threats, and select countermeasures to contain the expected losses within an acceptable threshold of risk. As with information assets, rings of protection—with differing types of security in each ring—are a good strategy.

Let's first look at identification of assets—the things that must be protected. Four physical asset classes are identified here:

- ◆ **Facility**—Building, rooms, workspace, backup storage area, and so on
- ◆ **Support**—Air conditioning, fire systems, electricity, communications, water, fuel supplies, and so on
- ◆ **Physical and components**—Hardware, including servers, printers, storage units, laptops, and workstations; desks; chairs; containers; and similar objects
- ◆ **Supplies and materials**—Disks and other removable media, paper supplies, waste material, and so on

Table 10.1, reproduced from a Royal Canadian Mounted Police (RCMP) presentation, indicates a number of specific protective measures, and it indicates to which of the asset classes identified previously each applies. The columns are labeled Facility; Support; and Supplies, Materials, and Components (the last column combines physical components and supplies and materials because applicability of the countermeasures is the same for both asset classes).

TABLE 10.1

**PHYSICAL AND ENVIRONMENTAL SECURITY—
PREVENTIVE TECHNIQUES/COUNTERMEASURES**

	<i>Facility</i>	<i>Support</i>	<i>Supplies, Materials, and Components</i>
Site location	X	X	X
Perimeter security	X	X	X
Construction standards	X	X	
Security containers			X
Drainage water detection	X	X	X
Access control procedures	X	X	X
Doors	X	X	X
Locks, keys, cards	X	X	X
Recognition badges	X	X	X
Access control logs	X	X	X

	<i>Facility</i>	<i>Support</i>	<i>Supplies, Materials, and Components</i>
Maintenance logs		X	
Transportation	X		
Fire protection	X	X	X
Offsite facilities	X	X	X
Waste disposal			X

Classification of assets also can serve a further purpose, one that is beyond the scope of this domain. That purpose is the risk assessment process, which helps determine which assets require how much protection from the threats and vulnerabilities explored here. Therefore, keep these broad distinctions in mind when learning the specifics of physical security, but remember that the classifications are an aide to learning and fulfilling objectives and do not specify an equal division of duty. Additional general topics in this chapter will refer to these classifications, whereas some specific topics will drill down further on a specific classification.

VULNERABILITIES

Understand some of the most common vulnerabilities and how they affect different asset classes differently.

Vulnerabilities affect assets. A common list of types of vulnerabilities is “destruction, disclosure, removal, and interruption.” At this level of abstraction, disclosure makes little sense—these are physical assets. Information assets (including things such as plans for physical assets like buildings or surveillance systems) can be disclosed inappropriately; physical assets themselves cannot.

The primary vulnerabilities of the classes identified here are

◆ Facility

Destruction:

- Accidental (fire, flood, earthquake, wind, snow, construction faults)
- Deliberate (vandalism, sabotage, arson, terrorism)

536 Part I EXAM PREPARATION

◆ Support

Destruction:

- Accidental (fire, flood, earthquake, wind, snow, construction faults)
- Deliberate (vandalism, sabotage, arson, terrorism)

Removal:

- Accidental (equipment failure, public utility outage, fire, flood, earthquake, wind, snow, construction faults)
- Deliberate (sabotage, vandalism, arson, terrorism)

Interruption:

- Accidental and deliberate are same as previous lists.

◆ Supplies, Material, and Furniture

Destruction:

- Accidental (fire, flood, earthquake, wind, snow, and so on)
- Deliberate (arson, vandalism)

Removal or Disclosure:

- Accidental (carelessness)
- Deliberate (theft)

Interruption:

- Accidental (fire, flood, and so on)
- Deliberate (sabotage, arson, vandalism, terrorism)

There are more elaborate systems for classifying assets, threats, vulnerabilities, and exposures. Such schemes are needed when performing a threat and risk assessment, but they would complicate things unnecessarily in this chapter.

THEFT AS THE MOST LIKELY PHYSICAL SECURITY ISSUE

Although it is not discussed as a specific topic, theft is one of the most likely security issues to affect an organization.

Employees are considered the most likely perpetrators because they have authorized access to sensitive information and valuable physical assets. This is simplistic; for example, laptop and hand-held computers can be more at risk to unauthorized people simply because they are small, portable, and valuable and are outside access control perimeters much of the time. Also, almost anyone is “authorized” when equipment is open to the public—for example, computers in a library.

In general, and just as simplistic as the comment about exposure to authorized personnel, theft is controlled by the following:

- Authorizing (or hiring) trustworthy people
- Maintaining a corporate culture in which honesty is expected and normal
- Motivating people by good work environments and competitive remuneration
- Minimizing opportunities that would allow the easy theft of assets

It is at the same time simple and difficult to be specific about theft because any physical or information item could be subject to this risk. General statements are difficult and probably misleading. Specific measures (such as computer cases that lock to protect valuable cards or chips, or cables to attach computers to something hard to move) are easy but not necessarily widely applicable. Lighting areas containing assets or located near valuable assets can go a long way toward making an asset less desirable as a target for theft. Lighting as a deterrent also involves sensors, and outside lights in approach paths.

Many physical security measures contributed to the control of theft (control of access and opportunity), or at least identification of the thief (through surveillance, logs, and other measures).

SELECTING, DESIGNING, CONSTRUCTING, AND MAINTAINING A SECURE SITE

Know the elements involved in choosing, designing, constructing, and maintaining a secure site. Elements include

- **Site location and construction**
- **Physical access controls**
- **Power**
- **Environmental controls**
- **Water exposure problems**
- **Fire protection and prevention**

Here is the crux of the issue: Your ability to physically secure assets depends on your ability to physically secure the site as well as the data center. A number of elements contribute to vulnerabilities, applicable threats, and the countermeasures that can be taken to mitigate them. In evaluating each site, not everything will be as easy to control. In studying the principles outlined here, you must realize that, although some risks can be eliminated or reduced due to proper site selection and facilities construction, we are rarely given that opportunity—and even these ideal conditions will vanish as time changes them and new threats appear.

The study of site selection, construction, and maintenance can best be understood within the framework of the controls available to mitigate the vulnerabilities previously described. These controls are roughly divided into the following:

- ◆ Site location and construction
- ◆ Physical access controls
- ◆ Power issues and controls
- ◆ Environmental controls
- ◆ Water exposure problems and controls

Site Location and Construction

Where the building is and how it is built are measures that significantly affect the level of vulnerability to threats and how well they can be mitigated. If the security team has the luxury of considering the location and construction of a new building (or remodeling a building), the following need to be considered:

- ◆ **Vulnerability to crime, riots, and demonstrations**—Is the location in a high-crime area of a city? Are you planning to construct a nuclear power plant on the San Andreas Fault? Will your staff be comfortable and safe leaving after hours in a dimly lit warehousing district? Is an unlit parking lot hazardous to night staff? These and similar questions need to be asked. Access considerations such as long, straight lanes or roads (where a truck could build momentum to crash through a wall) can be relevant if terrorism is a consideration. Nearby police and fire stations also could be factors.
- ◆ **Adjacent buildings and businesses**—Does a nearby business attract types of attention you don't want directed toward your information systems facility? If there is an adjacent building, can someone get from it into yours and, if so, is its security as strong as your own? A weak point in many homes is an attached garage; it often is less secure than the house and provides cover and tools for an intruder to spend time getting into the house proper. The same principle applies to adjacent buildings.
- ◆ **Emergency support response**—This already has been referred to: Nearness of fire stations affects how great your fire risk is, for example.
- ◆ **Vulnerability to natural disasters**—Is the proposed location susceptible to earthquake, tornadoes, or hurricanes? Is it located below a dam? Is it in an approach path to an airport? All these and other factors need to be considered. Government statistics from groups such as the United States National Weather Bureau help in assessing such threats as weather and other natural phenomena. Flood plain maps, earthquake risk maps, and similar data are available as well. It might be wise to consult an engineer or architect if more detailed information is needed; unless the security person is also qualified in such areas, risks can be missed.

- ◆ **General building construction**—Building construction is a major topic in itself. Obvious issues that should be considered include
 - Can the structure withstand hurricane-force winds (if relevant)?
 - Is it earthquake-resistant?
 - How many doors does it have, and how strong are they?
 - Will the roof withstand expected snow loads?
- ◆ **Computer room considerations**—In 1969, a computer center at Sir George Williams University in Montreal (now Concordia University), which was on display behind large glass windows as was popular then, was destroyed by gasoline bombs during a student demonstration. The computer center (whether a mainframe installation, a network server, or a server farm) should be a protected (point security) area within the building.

Even in an existing building, a computer center can be made fairly secure with little change to the existing structure. Full-height fireproof walls (to close off access through a false ceiling and some fire exposure) often are not especially expensive. (See the “Water Exposure Problems” section later in this chapter for more information on fire prevention.) Shatterproof glass and good locks on doors are other fairly inexpensive preventive security measures.

If alternatives are available, the location of a new building and its construction should be considered in the risk analysis and control program. Even if a new building is out of the question, secure areas for information systems within existing buildings usually can be added at a reasonable cost.

Physical Access Controls

Physical access control is essentially a perimeter control. You need to understand the following issues related to physical access controls:

- ◆ Perimeter control
- ◆ Access versus security tradeoff

- ◆ Response
- ◆ Doors
- ◆ Keys, including card systems and other tokens, and window construction

Some areas, such as computer rooms and rooms where computer media, servers, or data are stored, should have restricted access. Such areas need to be identified and marked. “No Admittance” signs do deter many people, and signs are very inexpensive. For greater exposures and potential losses, more expensive measures might be appropriate. These could include mantraps (entrances that permit only one person at a time to pass and that usually can be locked to trap an intruder) and various gates, fences, and detection sensors. Specialized knowledge is needed to design such perimeter controls to allow for issues such as emergency escape (fire exits, for example). Both active and passive measures should be considered.

Active Physical Access Controls

More active measures require people or, in some cases, expensive automated measures such as a computer-controlled card-access system. The people could be guards or receptionists. In either case, persons wanting to enter restricted areas should be preauthorized or accompanied by someone who is authorized. Some system of identification cards or badges normally is required to identify authorized personnel, unless the company is so small that everyone knows everyone else. Disaster planning should consider personnel access as well; many security procedures break down for janitors and are completely useless in stressful situations requiring access by emergency response personnel.

One thing that guards or receptionists should do is to ensure that access logs are maintained. Anyone (authorized or not) entering a restricted area should log in and out. The use of closed circuit TV (CCTV) as an “area” control might be appropriate to detect unwanted inhabitants. Use of force, including deadly force, might be appropriate—but it must be thought out carefully because many legal issues are bound to arise.

What has been discussed so far are essentially preventive and some detective controls. Reactive or corrective controls also should be included; a log of who is inside and when they are inside is not much good unless someone reviews it from time to time.

WARNING

Access Versus Security In essence, security is a trade-off when compared to access. More access implies a lower level of security. Each organization must choose the level of exposure consistent with its desired ease of access. A caution is that if security measures interfere seriously with what is perceived to be “normal” operational activities, people often will defeat the security measure. An example, seen frequently, is a door propped open because the automatic lock interferes with access to something as mundane as a vending machine.

Procedures defining what receptionists should do if someone unauthorized is discovered should be defined as well.

Passive Controls

Passive measures of access control include doors and locks. The doors should be of solid construction; making them fireproof can be a good idea because they then will also be solidly constructed. Reasonably secure locks are fairly inexpensive, but often are not provided unless specifically requested. Alarms to indicate that doors are open might be reasonable measures, if someone is monitoring the alarms.

There are many types of locks. Combination locks as well as keyed locks are available in various secure levels. Combination locks are more difficult to open in normal use, but combinations can be changed more easily than keyed locks can be re-keyed, and it is easier to keep track of combinations than of a rack of keys. Also, even though people can forget combinations, they cannot be lost as keys can. (Of course, if combinations are written down rather than memorized, the paper with the combination can be lost!)

For situations in which more sophisticated control is appropriate, more expensive lock systems—including remote control, magnetic locking mechanisms, and such—might be advisable. Such systems often are combined with access cards (“smart” or not) or other tokens, with or without biometric elements (fingerprints, pictures, facial bone structure, retina patterns, hand geometry and so forth). *Dumb* cards usually have a magnetic stripe that stores roughly 80 bytes, enough for basic personal information and some authorization codes. *Smart* cards contain processors and can include several kilobytes of information, enough for considerable biometric data and detailed records of what the token holder is authorized to do or has done. Smart cards can include enough processing power on the card to deal with encrypted communication to the control site, a major leap forward in security because many types of attacks become infeasible with encryption technology.

Normally, a computerized control system keeps logs of entry and exit, and this provides an access log and audit logs without the need to keep track of paper.

The two major considerations of what type of token to use are cost and safety.

The safety issue arises when deciding on failure modes (what happens to the doors when the system fails for any reason):

- ◆ **Fail-open**—Means that a power outage or computer crash can defeat the lock system. So, for any real security, serious Uninterruptible Power Supply (UPS) capability is essential (particularly if the locks are magnetic and require significant power to hold doors closed). This can increase costs substantially.
- ◆ **Fail-closed**—Appropriate in some high-security applications and for specialized cases, such as prisons. Fail-closed means there will be no exit from a fire and thus usually contravenes many laws and regulations.

Costs depend on the type of card and the type of system. Non-smart cards are not reusable and are cheap. Smart cards cost in the range of \$2 (in large volume) to \$7 or \$8; specialized cards can be very expensive. Smart cards usually are reusable, which helps somewhat to mitigate costs. Generally, a system involving smart cards implies significant computer and communications capability (between sensors and central database processors) and can be expected to be relatively costly. Where biometrics are involved, the sensors that read fingerprints, cameras that “look at” faces or retinas, and other biometric sensors also are more costly than simple magnetic strip readers.

Systems involving biometrics have other issues, such as reliability and errors. In 2002, biometric sensors are an evolving technology; false positive or false negative errors usually are in the range of 0.01%–1.0% for the better-developed technologies like fingerprint readers. This sounds good and is acceptable in some cases (usually when the traffic volume is relatively low). However, a 1% false positive rate in an airport with 100,000 passengers daily means 1,000 people are flagged incorrectly every day, or 3 people per jumbo jet—and there are many airports with far more than 100,000 passengers daily. Some technologies, such as face recognition, have error rates closer to 5%.

As is repeated many times in this chapter, risk assessment (related here to perimeter access control) should identify the threats, vulnerabilities, exposures, and an acceptable loss; smart card or other token systems might be cost-justified.

Attention should be paid to windows as well. If windows allow individuals to look inside an area where sensitive data is handled, the sensitive data should not be visible from a nonsecure (outside the area) side of a window. Windows might need special construction to resist attack or even weapons fire. (Films can be applied to standard windows that provide considerable resistance to even small arms fire, for relatively little expense.) If sophisticated eavesdropping attempts are part of the threat profile, windows might need to be opaque to various wavelengths (infrared, for example) or might need special mounting and materials to prevent extremely sophisticated listening techniques, such as the use of lasers that can measure the vibration of glass due to conversations inside and therefore determine what has been said. Window frames need to be designed and installed in keeping with such special considerations as they apply.

Power

Computers need electrical power to work. This area is a technical one in which detailed examinations require specific technical training, and an expert should be involved in the design process.

The first level of expert is the manufacturer of the computer(s). Pay attention to what type of power the maker says should be supplied.

Most computers are sensitive to *dirty* power (a power supply that has significant voltage variations, interference, and similar variances from what should be expected). A consideration for microcomputers, for example, could be other office equipment on the same power line. Some electric typewriters generate a fairly powerful short surge when the carriage return is engaged. Such a surge in computer equipment attached to the same power line is not good, so protection is needed. The first rule of computer power usually is “isolation”—the computer should be on a different line than other office equipment. This rule applies to personal computers as well as to mainframes. (Practically, manufacturers have made personal computers relatively insensitive to this sort of power fluctuation; otherwise, no one could use them at home.)

Power supply conditions should be monitored. Many automatic devices are available that will keep a record of usage and similar items. From a security perspective, you should consider the building’s electrical room as well; penetration here could stop the computer as surely as penetration into the computer room itself.

Relatively cheap surge protectors and filters can protect computers from most dirty power problems; a power supply monitor allows a designer to know what sort of filtering is necessary.

Power Issues: Spikes, Surges, and Brownouts

Computer equipment is vulnerable to many things in the power supply. The most common risks are as follows:

- ◆ **Brownouts or total power loss**—The voltage drops, or power is lost entirely in this case. Some disk drives and other motors can be very sensitive to low voltage. Some processor and memory chips can have their lifetimes significantly reduced in an environment of significant brownouts.
- ◆ **Spikes and surges**—These result when lightning hits outside power lines, or in some other circumstances, when a sudden spike of high voltage appears on the power or telephone lines. Computer equipment and modems connected to telephone lines are very sensitive to high voltage spikes.

Surges are common on lines with electric motors attached. The voltage drops a bit when a motor starts and then surges a bit when it stops.
- ◆ **Static**—Particularly in cold climates, people generate static electricity when moving around. If the humidity is low, sparks are common, and a spark can ruin a computer chip or scramble data in a flash memory chip. At a minimum, data currently being processed in the computer can be corrupted.

Minimizing Power Problems

One way to minimize problems with power is to install a UPS. The level of UPS needed can range from batteries that will support the system for a few seconds so that it can fail soft (that is, shut itself down controllably), to elaborate systems including backup generators for systems that must continue to function regardless of the failure (air traffic control or hospital systems, for example). UPSs can cost \$100–\$200 (for a small system that will run a personal computer for long enough to finish copying files onto a disk) to \$100,000 or more (for elaborate battery systems with automatic backup generators).

When the computer *must* keep running, or when it is convenient to allow a soft shutdown, some self-contained power supply units can save a lot of trouble—they will detect the eventual loss of power due to their battery exhaustion and shut down the computer in an orderly manner. This prevents the damage that can result when a system fails or crashes due to low power.

Protection from many power supply problems is fairly inexpensive. In many areas, good solid, even power is available. Few surges and spikes and few incidences of power outage occur. Interference is thus unlikely if the computer has its own power line. Weather, however, can be the major risk. For some, unplugging modems from phone lines and all equipment from power lines during lightning storms eliminates surge risks due to these storms. (If lightning gets into the inside power lines, you have more problems than just a fried computer!) Most companies, however, want to ensure that appropriate surge protection and backup power are available so that work can continue regardless of the weather. The cost to do so is directly related to the amount of computing that is critical to business operation and is inversely related to the nature of the power supply and the interruptions caused by weather and other events.

Static is minimized by controlling humidity. Antistatic mats under chairs and machines and antistatic carpeting are advisable in areas prone to low humidity, such as cities in very cold climates during the winter. Antistatic sprays (marketed to stop clothes from clinging) can also help around computers.

Many less expensive UPS systems provide only a few minutes' supply—enough time to allow you to save files and shut down softly. The short-term supply necessary for this might be available for as little as \$150–\$200. A UPS that provides more long-term protection starts in the cost range of \$500 and goes up quickly, depending on the wattage capacity required. Many microcomputer installations do not need the extended capacity of these more expensive systems, but a lot of time, trouble, and lost work might be avoided with one of the less-expensive units. UPSs might or might not also include filtering of interference, spikes, surges, as well as backup power. You should, however, fully qualify what a unit does before purchasing and putting it into operation.

More elaborate UPSs include such capabilities as power-generating facilities that are intended to start up to maintain continuous power after a battery UPS discharges. It is critical that such systems be tested frequently, and such testing must include extreme conditions (for example, a diesel generator outside a building might start perfectly in August but fail to start at -40°F in January).

Manufacturers are aware of the sensitivity of their equipment. The better microcomputers typically have some built-in protection and occasionally even some very short-term power backup (to avoid problems with millisecond blips). No-name units often skip these features and might not be as good a buy as they seem for this reason.

You should not buy a computer that is not Underwriter's Laboratory (U.S.) or Canadian Standards Association (CSA) approved. *Never* buy a power supply that is not Underwriter's Laboratory (UL) or CSA approved. UL or CSA approval relates to safety features, not to performance; non-approved equipment can be a hazard, violate insurance policies, and be unlawful.

Environmental Controls: Air Conditioning, Humidity, and Temperature

Most large computers require special air conditioning to continue to function properly. This can extend to smaller systems as well; for example, it is not unusual to see someone begin to experience copier problems when a copier is enclosed in an improperly air-conditioned room.

Again, the manufacturer is the first source for expert advice. The maker should specify cooling requirements, and the user should heed the specifications.

As with power, the air conditioning for a computer should be for the computer only. It makes no sense to try to share the load with other, unrelated areas and risk expensive computer hardware.

Air conditioning units require supplies of air and often water, and they generally produce water from condensation. Fire prevention includes making sure the fire won't find a ready entry to the computer through the air conditioner. Water supplies must be controlled to ensure that busted pipes won't destroy the hardware.

As electricity-consuming equipment, the air conditioning needs its own power—separate from the computer. Often, a second cooling unit is appropriate to ensure that if one fails, all cooling is not lost and the system can continue to function.

Automatic humidity- and temperature-monitoring devices should be installed in climate-controlled computer rooms; the records should be examined regularly to ensure that the climate control is functioning properly.

As solid-state technology continues to improve, the amount of heat generated by computers and the resulting air conditioning need are decreasing. Most personal computers require no more “comfort” than people, and this is also true of some mainframes. In fact, the primary air conditioning problem found in offices with many microcomputers is uncomfortable people. A lot of computers collectively generate a considerable amount of heat, as do copiers and laser printers; offices not designed to handle the load can become very uncomfortable workplaces. This often affects productivity in a negative way; to help keep a happy company, the risk management team needs to consider cooling its people as well as the computers.

In a related issue, some laptops get quite warm when operating, especially if they are playing DVDs. Holding such a device on your lap for a prolonged period can cause physical problems from the heat.

Water Exposure Problems

Water exposure problems can be caused by something as simple as a window open during a rainstorm to something as wide-ranging (and outside an individual organization’s control) as a collapsed tunnel letting a river into most of downtown Chicago’s sub-basement system. A short list of common problems include

- ◆ **Flood**—Whether from weather or municipal facility problems
- ◆ **Basements**—Water from an upper floor problem tends to result in flooded basements
- ◆ **Roofs**—Leakage, burst drainpipes during heavy storms, and so on
- ◆ **Snow load problems**

- ◆ **Hurricane and other weather phenomena**
- ◆ **Sprinklers**
- ◆ **Air conditioning**—Often uses water as a coolant or heat transfer fluid

Careful attention to drainage can help with many of these problems, as can location of the computer room (obviously, all else equal, you shouldn't put the computer center in the basement of a building). Weather precautions vary depending on the local climate. Sprinklers do an excellent job of extinguishing fires, and if the water is clean, it might not seriously damage computer equipment. Sometimes, simply drying out a computer is sufficient, but sometimes specialized recovery techniques are needed. For events like the tunnel collapses in Chicago, insurance might be the only answer, unless location of the data center outside the risk area is feasible.

Fire Prevention and Protection

Fire prevention is not the same as fire protection. *Protection* refers to detecting fire and minimizing damage to people and equipment when it happens. *Prevention* is avoiding the problem in the first place and usually is less costly and more effective in minimizing damage.

Most jurisdictions have fire codes, which specify legal requirements for minimum fire prevention measures. Expert advice should be sought to ensure that the information systems activities conform to applicable fire code regulations.

Four elements of prevention are outlined in the following list:

- ◆ **Construction**—The materials used in a computer room should be as fireproof as practical. Combustible material (stacks of paper, for example) should not be stored in computer rooms, or indeed around any other electrical equipment. False ceilings should not be flammable. False ceilings and various parts of the heating, ventilation, and air conditioning (HVAC) system can provide “chimneys” to permit rapid spread of fires, and it might be advisable to close off such openings. Rugs, unless specially designed for the purpose, do not belong with computers (for reasons of static electricity as well as flammability).

Magnetic tapes and plastics such as CDs and DVDs are difficult to ignite when stored in containers, but they're also difficult to extinguish when ignited. Plus, they produce poisonous combustion products when they burn. If a media storage vault opens onto the computer room (a very common design, for excellent efficiency reasons), special attention is needed to minimize spread of a fire between the equipment and the media vault.

- ◆ **Training**—Fire regulations should be known and observed by all employees. Employees should be given training in fire prevention as well as in what to do when a fire does occur. The training should include instructions about exits, available extinguishing equipment, emergency power, and other shutoffs.
- ◆ **Testing**—Fire procedures should be tested periodically with fire drills. (This is normally required by local regulations. It's also a common-sense practice.) There is a risk here: Too few fire drills will not maintain familiarity with procedures, while too many will create a “boy who cried wolf” situation. In the case of a real fire, people might be slow to respond because they will think it is yet another drill.
- ◆ **No smoking policy**—For fire risk and other reasons, smoking should not be allowed around computers. This also applies to personal computers—the lifetimes of disks in environments with cigarette smoke might be very short indeed because the smoke particles can adhere to the media via static and other charges and cause read errors. Smoking also provides a source of ignition. Everyone probably has seen the worn tracks in carpets where cigarette smoking is common and ashes fall to the rug; a cigarette dropped into a waste paper box could cause a very destructive fire.

If prevention does not work, fire protection becomes the issue. The first thing is to detect the fire. Obviously, you want to detect it while it is still small and controllable.

Fire-detection systems are common and inexpensive. Ionization-type smoke detectors react quickly to the charged particles in smoke (remember what charged particles in cigarette smoke can do to oxide surfaces on disks). Photoelectric detectors, on the other hand, react to light blockage caused by smoke, and heat detectors react to the heat of a fire.

Combinations of these detectors can detect a fire very quickly, and often before there is a serious problem. Most local fire codes now require smoke detectors in residences and workplaces; the mass production of detectors has brought the costs down drastically. Effective smoke detection, including both ionization and photoelectric detectors, can be achieved for a small investment.

The first rule after a fire is detected (either by smoke, heat, or other means) is to *get the people out*. Fires can spread very quickly, more quickly than many people realize, and toxic gases are produced as well as heat and smoke. People are the most important asset and are difficult for an organization to replace, as well as having high intrinsic value. Only after all personnel are safe and accounted for is it appropriate to attempt to put out a fire, and then it should be done only after calling the fire department.

Many fire extinguishing systems are available. Portable fire extinguishers always should be available near any electrical equipment, including computers. These extinguishers must be examined periodically to ensure they remain useful. For computers, type ABC extinguishers are appropriate because combustible solids (class A), combustible liquids (class B), and electricity (class C) all are common in computer room fires. Get the people out *first*; then an attempt can be made to extinguish a small fire using portable or other extinguishers. The primary purpose of extinguishers is to ensure that an escape route can be cleared; the fire department always should be called and the people evacuated before any extinguishing attempts are undertaken.

Fixed systems include carbon dioxide extinguishers, with or without directing hoses. The entire computer room can be flooded with carbon dioxide to put out most fires by depriving them of oxygen to support combustion; with hoses, the gas can be directed at specific fire sites. Such systems are expensive and should not be automatic: They deprive people (such as computer operators) of oxygen, as well as depriving fires of oxygen. Installation of such systems is a job for professionals.

A fire-protection system that is safer for people and that extinguishes fires without irreparably damaging computer equipment uses Halon 1301 gas. This gas has the convenient property of smothering fires without being quickly fatal to people, so automatic systems can kill the fire while allowing people enough time to get out. Halon systems are installations requiring specialized expertise, so professionals should be engaged.

WARNING

Extinguishing Fires Any attempts to put out a fire must be done by people who have appropriate training. Choosing the wrong material can be hazardous to health. For example, attempting to put out an electrical fire with water can lead to electrocution. In the heat of the moment, this simple thing can be forgotten. Also, improper use of a fire extinguisher can spread a blaze rather than put it out. In addition, fires usually create toxic gases, especially fires involving plastics. Smoke inhalation of such toxic compounds kills more people than flame in many fires, sometimes including people who stay too long trying to put out a fire.

552 Part I EXAM PREPARATION

Halon systems also are expensive, as are tests of the system (a refill can cost more than \$1,000). Such elaborate fire systems probably are appropriate only in mainframe installations. (Halon 1301 and Halon 1211 are trademarks of chemical compounds, owned by Great Lakes Chemical Company Inc. The details of composition are not relevant in this text and are not public information in any case. Halon 1301 is not self-pressurizing and requires expensive pressure systems for a fire installation; Halon 1211 is self-pressurizing and can be put into a portable extinguisher, either alone or mixed with Halon 1211. Such portable extinguishers have been available as normal retail items; although this is no longer true they might still be in use.)

With the signing of the Montreal protocol in 1987, Canada, the United States, the European Community, and 23 other nations agreed to control the production and consumption of certain chlorofluorocarbon compounds (CFCs), including the Halon group. These ozone-depleting substances include some refrigerants and, relevant to this discussion, Halon 1211, Halon 1301, and Halon 2402. These Halons are used primarily in fire-extinguishing applications. The CFC compounds are implicated in the depletion of the ozone layer, a potentially serious global environmental problem.

The timetable for implementation of the Montreal protocols was advanced in 1992, and chlorofluorocarbon fire systems might not be a viable alternative for new, or even existing, installations. Halon systems are still used in special circumstances, but under severe regulation.

Regulations regarding the use of Halon vary, but typically include these recommendations:

- ◆ When planning fire protection for new installations, all alternative options (carbon dioxide, water, and so on) should be fully explored before deciding to use Halon.
- ◆ When Halon is used, full-discharge testing should be avoided in favor of alternative test procedures.

Alternative test procedures include a room pressurization test and the “puff test.” Standards for a room pressurization test are available from national fire-protection groups. The puff test involves putting lightweight caps over outlets and using air to ensure that piping is free from obstructions. The professional should ensure that correct procedures and relevant local regulations are known and followed.

It is a good idea to avoid water in computer room fires; automatic sprinkler systems normally use water. First, computer fires usually involve electricity, and water conducts electricity. Second, water is likely to seriously damage computer equipment and can do more damage than small fires. The fact remains, however, that water is an excellent way to extinguish fires—one reason it is used by fire departments. In the absence of electrical power, clean water should not damage computer systems, although they must be dried soon and carefully to avoid rust and corrosion problems.

A special problem often overlooked in using water to extinguish fires is how long the water has been sitting around in a building's pressurized system. Because fires often disrupt electrical power, building sprinkler systems often have separate water supplies, not dependent on outside electricity or piping. One way to do this is to have a reservoir somewhere high and separate pipes that are always filled. Another alternative is a separately powered pump system, usually located in a building's basement. (This is called a *wet standpipe* system; in a *dry standpipe* system, water is pumped into the building system from the outside by the fire response units.) Such reservoirs tend to be filled once and then checked for level periodically; it is rare to see checks for purity as well as level. The previous statement that water conducts electricity is not strictly true: *Distilled* water is not a conductor under normal circumstances. However, tap water, and especially water that has resided in a reservoir for an unknown time, is *not* distilled water. Some of it can be decidedly contaminated. The main lesson from a security perspective is that computer room fire protection should consider carbon dioxide or Halon, not the building's sprinkler system.

TAPE AND MEDIA LIBRARY RETENTION POLICIES

Understand issues and controls related to removable electronic media.

Computers work with data, and the data and information into which the data is processed generally need to be stored. This is the job of magnetic tape, disks, compact discs (CDs and DVDs), and other media. The list of media is long already, and grows daily.

(It also shrinks: Punched cards and 5 1/4" disks are no longer common.) Different media have different characteristics and different capacities. All media contain data, and the data on the media is just as valuable and just as sensitive in movable form as when being used by the computer. Removable media, by definition, also are at least somewhat portable. This presents a security and control risk. Usually it is recommended that there be a tape/media library for storage purposes.

Depending on the installation, the media library can range in size from a small cabinet to a rather large warehouse-size space.

Whatever the size, the media storage area should be

- ◆ **Restricted**—Storage areas need to be at least as carefully controlled as the area in which the data is used. Many computers are not especially portable, but removable media is. The equivalent of several books can fit onto CD that will fit easily into a shirt pocket. The equivalent of a large book will fit onto a memory stick, which can also be easily slipped into a pocket. (You might be familiar with memory sticks, which are used in digital cameras to store pictures and are about half the size of a stick of chewing gum.) If the book contains sensitive information, such as the corporate budget, careful protection is needed. All the access controls recommended for other restricted areas also are necessary in the media storage area.
- ◆ **Controlled**—Someone should have specific responsibility for keeping records of media entering the library and leaving it, and for conducting frequent inventory of the contents. Any discrepancies should be followed up immediately.
- ◆ **Locked**—This is an elementary issue, but it is frequently ignored. Some form of an automatic locking mechanism is preferable, so that carelessness cannot lead to a large exposure.
- ◆ **Protected from fire**—Media contain, as an acquired value, information that might be expensive or impossible to replace, and that might be valuable to others as well. The storage area should be separated from the rest of the computer resource and should have its own independent fire protection. This could be elaborate in a large installation or fairly simple in a small shop.

No general rules on fire and access protection are practical because media vary too much in their characteristics. Punched cards were flammable and had to be kept in humidity-controlled areas to prevent warping, which can cause feed jams. Magnetic tapes are sensitive to heat and burn fiercely but are not especially easy to ignite. Optical storage media are extremely long-lasting and are not fragile (but they have very high capacities and might need more careful protection because of the sheer volume of information they hold). Optical media are also plastic and thus a potential fire hazard. Flash memory systems such as memory sticks for cameras are tiny for the information they can hold and are not fragile, but they might need special measures because they are so small.

A basic rule is that any sensitive data should have at least two backups, and at least one should be stored in a different building separate from the others.

DOCUMENT (HARD-COPY) LIBRARIES

Understand issues relating to storage of paper.

Many considerations that apply to storage of media also apply to storage of paper documents. Security considerations are essentially the same, with the exception that the exposure due to unauthorized access is lower because information on paper is far less dense than on magnetic or optical media. Although the risk of a single document or a small number of documents being compromised is higher, the risk of loss of enough information to form a coherent overall picture is much lower. (Of course, some single documents can be highly critical, just as some data files can be unusually sensitive.)

In terms of physical storage, paper is more resistant to heat from fires than are magnetic and optical media. Offsetting this, paper is much easier to ignite, and such fire suppressants as water will damage paper seriously when they might not significantly affect plastics. Therefore, physical storage for paper documents needs to be

- ◆ Larger in volume than for magnetic media
- ◆ Protected from water damage more carefully
- ◆ Treated as a fuel repository and kept well separated from more sensitive media

The following is a useful checklist (adapted from *Disaster Planning for Government of Alberta Records*):

- ◆ Keep passages unobstructed.
- ◆ Do not store records on the floor.
- ◆ Do not leave original documents on desks overnight.
- ◆ Store cellulose-based nitrate films separately, and treat them as flammable and hazardous goods.
- ◆ Do not pack files too tightly (water can cause swelling and burst packaging).
- ◆ Set materials back slightly from shelf edges to lessen vertical fire propagation.
- ◆ Avoid basement storage.
- ◆ Check areas where condensation can be a problem (pipes, windows, and so on).
- ◆ Install shelving at least 12" from outside walls and 2" from inside walls, and place bottom shelves at least 4" above the floor.
- ◆ Store more valuable material on upper shelves and upper floors.
- ◆ Avoid carpeting in storage areas.

WASTE DISPOSAL

Know the most common issues related to disposal or erasure of data.

One of the classic computer crimes reported in the literature involved a person gaining accounts and passwords to get into a computer system, and instructions on how to compromise it, by going through a telephone company's waste bins. (This often is called *dumpster diving*.) Similar incidents have involved statistical and taxation data. The security and control principle here is that discarded listings, media, and anything else containing data or information remain sensitive (if they were in the first place). Control on disposal is necessary.

Classified wastes should be

- ◆ Stored in separate containers
- ◆ Collected frequently, by security-cleared personnel
- ◆ Retained in a secure area
- ◆ Destroyed by cleared personnel, using an approved and effective method (shredding, incineration, and so on)

Note that the cleaning staff must be cleared or kept out of areas containing sensitive assets.

Some points should be kept in mind here:

- ◆ Most personal computer operating systems do not actually *erase* data files when the operator says “erase” or “delete”; they set a flag indicating the file is “deleted.” The flag can be reset, and fragments of data might still exist. (Some of application software also does not necessarily destroy data when you delete it: For example, many database products don’t delete items until the database is *packed*.) In fact, programs exist specifically for the purpose of recovering deleted files. *Degaussing* is needed to ensure the erasure (a degausser generates a strong, varying magnetic field that randomizes the magnetic bits used to store data).

Note that formatting a disk on a personal computer might not destroy data (this depends on the operating system and hardware manufacturer). Overwriting, degaussing, or physical destruction is necessary.

- ◆ Data stored on most commonly available optical media (such as CD-ROM and DVD) cannot be erased; the medium must be destroyed thoroughly. However, read/write optical systems are becoming common. Read/write optical media (CD-RW and some DVD) are erasable. WORM (write once, read many) systems, including CD-R and DVD, act like read/write but actually simply use the enormous capacity of an optical disc to store multiple copies of data, one for each version. WORM has advantages where a record of historical changes is necessary; the key here is that the data cannot be erased.

- ◆ Core dumps generated during program development (or sometimes when a program fails during operation) are sensitive waste. They contain a great deal of information that can be read by trained personnel, sometimes from areas outside the specific program's authorized accesses. Listings must be controlled as classified waste.
- ◆ Some kinds of computer memory stay "live" for a long time (up to years) even with the power turned off. An unauthorized user turning on the machine might get access to sensitive information unless the memory is actually written over with 0s, or some similar destruction method is used.
- ◆ As mentioned previously, data on magnetic media usually is nonvolatile. If you put a customer list or proprietary information on a fixed disk and then sell or trade in the computer, format the disk before it leaves your premises.

Degaussing is a coined word relating to removing magnetism (a *gauss* is a measure of the strength of a magnetic field). Disks should not lean against a telephone; I also could add, "Don't put a disk on top of a television or audio speaker." A *degausser* is something with a strong magnetic field, preferably a moving field, which is not the same as the fields that write to magnetic storage media. (In computer terms, properly degaussing *removes* magnetism, and the discussion here is merely of *changing*.) The magnet that rings a telephone bell, moves the cone in a speaker, or controls the picture tube in a television induces a magnetic field that is not at all like data on a disk. Magnetic media are designed to capture and retain imposed fields; the media don't care what the patterns are. The computer decidedly *does* care. Most firms that deal with magnetic tape have bulk tape erasers (it's much faster than doing it with a tape drive, and tape drives have more valuable uses). A recent edition of a commercial catalog lists a "Magnetic Bulk Tape/Floppy Disk Eraser" for \$39.95. If sensitive material is stored on magnetic media, a degausser can be very cheap insurance, if it is used regularly.

Security personnel should recall that data stored on *optical* media has a very different, nonmagnetic means of recording, and magnetic fields (and degaussers) are irrelevant. Except for read/write optical media, optical discs *cannot be erased*. Even considerable physical damage might not destroy the data. One favorite demonstration of optical disk sales people has been to pour coffee, cream, or some such liquid onto a disc and then wipe it off and proceed to read it.

(This works better with black coffee; you need to use soap and water to remove sugar and other sticky stuff.) To dispose of an optical disc, physical destruction is necessary—breaking it into pieces or melting it works best.

When disposing of classified data, more stringent rules might be necessary. File wipe programs exist that actually overwrite media, rather than merely deleting the contents or directory entry. Although some file wipe software uses particular patterns of bits to ensure the maximum chance of overwriting everything, there are issues of physical play in read/write heads and of remanance in the media. Advances in technology have made it possible to read nearly any magnetic pattern that ever was imposed onto magnetic media; even a file wipe might not be sufficient for classified material. Physical destruction of media might be required.

In the special case of nonremovable media that need repairs or are being discarded, consideration must be given to the risk of advanced techniques being used to read the waste. (Of course, advanced techniques are unnecessary if a disk being repaired has not been wiped, as noted previously.) It is common in high-sensitivity situations to destroy any media, removable or not, that must leave a high-security area for any reason.

More information about dealing with classified material is found in the “Government of Canada Industrial Security Manual,” and in Department of Defense Guide DOD 5220.22-M on sanitizing media.

NOTE

Offsite Storage Data (or whatever) stored offsite (somewhere outside the normal computer center) must have a level of security and control at least as good as the computer center has. Extremely tight security in the computer center does little good if backup copies of the same data and information are unsupervised in a warehouse without adequate fire or access control. The same considerations apply while media are being transported.

PHYSICAL INTRUSION DETECTION

Describe physical intrusion detection methodologies and products.

For this section, physical intrusion detection is presented in the form of a fairly long table. A key point to remember: Do not be blinded by technology; sensors and detection must be guided by intelligent risk assessment and must be part of a complete strategy.

However good a perimeter security mechanism is, sooner or later it will be defeated. When this occurs, there should be a method for detecting the intrusion. The key to a good defense is a defined monitoring and response mechanism.

Many kinds of sensors and other detection mechanisms are available. Table 10.2 lists some of these mechanisms and some characteristics and issues related to them.

TABLE 10.2**SENSORS AND OTHER DETECTION MECHANISMS**

<i>Sensor</i>	<i>Description</i>	<i>Issues</i>
Motion detector	Can use infrared light beams, lasers at any wavelength, microwaves, or other means to detect motion in an area where there should not be motion. Sensor units that broadcast a signal can be about the size of a pack of cigarettes, and the receivers can be small.	Can be installed to detect an approach to a perimeter or presence inside a controlled area. Can be very inexpensive or more expensive and very sophisticated. More sophisticated installations require power to operate a central processor; less sophisticated installations usually run on batteries (which, of course, must be checked and replaced periodically). Requires some sort of response system to determine what caused the sensor to trigger. Normally installed out of sight, although light-based units require a line-of-sight to the area to be monitored. Susceptible to triggering from natural events such as wind.
Heat detector	Measures increased temperature from a heat source—fire, human or animal body, or other source. Sensor unit can be very small (millimeters) if at the end of something like a lens attached to an optical fiber.	Use and considerations as for motion detectors. Does not require line-of-sight and does not react to wind or most natural events. Can also detect fires, sometimes before there is an actual flame. Can react to small animals if sensitivity is set high.
Vibration sensor	Measures vibrations caused by events such as glass breakage, collision of a vehicle with a wall, footsteps, or other noises. Can react to noise, broken foil on a window, and magnetic or mechanical switches.	Use and considerations as for motion detectors. Often visible in the form of tape on windows. Can be a very sophisticated system such as a laser measuring displacement. It's susceptible to triggering from wind and other natural phenomena.
Capacitance detectors	Measure the change in capacitance caused when an animal or a human approaches the sensor.	Usually installed on things like fences. Susceptible to false alarms from wild animals (including raccoons and dogs even in cities). Sensor units do not have to be close to the point of interest.
Magnetic sensors	Measure changes in magnetic fields caused by the presence of a conductor. Well-known examples include gates and wands in airport security screening.	Typically react to conductors, including innocuous items such as keys and coins. Will not detect things such as plastic or nylon knives or explosives.
Sniffers	The best of this class of sensor is a trained dog. Canaries are used in some situations as well. Technological solutions include some type of device that collects air and performs tests to determine the presence of items of interest. This form of technology is fairly expensive and must be considered a developing technology.	Response time of machine-based sensors can be slow. Dogs and machines react only to those things they have been trained or built to detect. Animals are susceptible to fatigue fairly quickly and require trained handlers and controlled environments. Nonliving sensors do not fatigue and can be very useful for applications where response time is not critical.

<i>Sensor</i>	<i>Description</i>	<i>Issues</i>
X-rays and other see-through devices	Can require significant power supplies. X-ray and other radiation technologies are common. Time-modulated ultra-wideband is a relatively new spread-spectrum technology that allows handheld devices to use low-power radio waves as a radar to see through walls, clothing, and so on.	These sensors can provide images of the contents of sealed containers. The most common uses are in baggage screening and police work (to see what's going on inside a building). Some radiation devices produce sufficient energy to affect things such as unexposed camera film. In all cases the view can be difficult to interpret, and operators of such scanning tools frequently misidentify items in orientations different from "normal" (for example, a knife looks very different from the hilt end than from the side).
Cameras	Can range from simple CCD (charge-coupled device, a small and low-power-drain type of imaging chip) units providing a video feed to a central monitoring station to pointable cameras. Can be sensitive to infrared, ultraviolet, or other invisible frequencies. Some variations can see in complete darkness, usually by infrared. Individual units can be small, and there might be nothing other than a lens (perhaps 1mm in diameter) attached to an optical fiber, with the actual sensor remotely located.	Requires monitoring or some form of recording. More sophisticated systems require power; CCDs and similar devices need only small batteries. Cameras are so small and inexpensive that they are becoming ubiquitous; people are rarely out of view of a camera in many workplaces and public spaces. There are significant issues of privacy in public areas. Monitoring can be a problem because most of the time it is a very boring job.

Table 10.2 is not an exhaustive list. Perhaps the biggest problem with such devices is that often there's a tendency to install sophisticated sensors without appropriate threat and risk assessment, real-time monitoring capability, or well-defined response procedures. As with other detective controls, if an appropriate response plan doesn't exist, there is not much point in installing the latest technology.

IN THE FIELD

PHYSICAL ATTACK PARAMETERS

Several observations regarding site selection and building construction have been made in this chapter. In situations of national security or where terrorism might be a factor, careful attention must be given to measures that will lessen vulnerability to physical attacks.

Many sources (for example, Van Nostrand Reinhold's *Computer Security Risk Management* and *RCMP Security Information Publications # 3*) provide information about typical times for various methods of physical penetration and using various tools.

continues

continued

The tactic of putting power poles in the back of a truck and backing it at a high speed into a wall has been mentioned (as has the counter of minimizing long, straight lanes and roadways). Similarly, chain link fences can be penetrated with minimal damage to vehicles but can be strengthened substantially simply by attaching a cable to back up the links.

Typical times and other considerations vary greatly among these lists, frequently for the same attack using the same tools. Also, the introduction of weapons into a situation materially changes things. Nevertheless, such lists can serve as a guide to physical security measures related to site construction and selection. Perhaps the single most important message from the lists is this:

Multiple rings of protection, with different preventive measures requiring different tools for penetration at each barrier, can slow an attack significantly, allowing response teams to arrive.

CASE STUDY: BLOWING UP SECURITY—THE CASE OF THE BALLOON

ESSENCE OF THE CASE

Our systems admin was working late and left the data center to visit the food machines in the cafeteria. Upon his return, he found himself locked out of the center. He had left his access card within the data center. Like many facilities, his center required the insertion of a security card in order to enter. A valid card triggers a release mechanism and the door opens. Anyone with a card can enter. To leave this particular system, however, is easier. Motion detectors on the inside of the data center detect someone moving toward the exit and open the doors.

SCENARIO

We can learn a great deal about physical security by studying the vulnerabilities discovered by others. Often these penetrations are the result of real-world attacks, but sometimes they result from accidental discovery. In this case, the perpetrator was the systems administrator. He meant no harm; he merely had left his access badge within the data center and needed to return. It was after hours, and no one else was around. This story comes from a discussion found recently on the Internet. The names of the participants and the company are not revealed, in case the vulnerability has not been addressed.

CASE STUDY: BLOWING UP SECURITY—THE CASE OF THE BALLOON

Our resourceful admin recalled that earlier in the day a birthday had been celebrated in the reception area. He returned to the area and found the penetration tool he desired—a balloon that hadn't been blown up.

He returned to the data center and laid down facing the entrance doors. He pushed the balloon under the door, leaving the mouth of the balloon on his side of the door. Holding the neck of the balloon between thumb and forefinger, and placing his lips over the mouth of the balloon, he began to blow. As he blew, the balloon grew in size—on the inside of the data center. You can imagine the rest. He released the balloon and jumped up. The balloon flew around the immediate inside of the data center and triggered the motion detector. The doors opened, and our administrator was able to enter the facility and continue his work.

ANALYSIS

Even though this is a humorous account and no harm was done, it points out the need to review physical security devices and look for even the most bizarre vulnerabilities we might find. If the doors had been flush with the floor, instead of providing a handy gap, the penetration would not have occurred. If the motion detectors were tuned (if capable) to respond to a range of motion not within the purview of a rapidly decompressing balloon, the penetration would not have happened. If, of course, motion detection was not used to open the doors from within, the penetration would not have occurred.

CHAPTER SUMMARY

Physical security refers to the provision of a safe environment for information processing activities and to the use of the environment to control the behavior of personnel. This chapter has addressed these issues by discussing your need to

- ◆ Know the elements involved in choosing, designing, and configuring a secure site.
- ◆ Know how to secure a facility against unauthorized access and theft of equipment and information.
- ◆ Know environmental and safety measures needed to protect people, the facility, and its resources.

KEY TERMS

- Area control
- Clearing
- Core dump
- Degauss
- Degausser
- Dirty power
- Dry standpipe system

continues

CHAPTER SUMMARY *continued*

- Dumpster diving
- Erasure
- Escort
- Fail-closed
- Fail-open
- Lock and key protection system
- Magnetic flux
- Magnetic remanence
- Media
- Memory stick
- Open storage
- Overwrite
- Perimeter control
- Physical control space
- PIDAS (perimeter intrusion detection and assessment system)
- Purging
- Restricted area
- Sanitization
- Security area
- Security perimeter
- Survivability
- Threat profile
- Wet standpipe system
- WORM (write once read many)

It is now up to you to formulate in your own mind and words an approach to physical security that specifically addresses the needs of your facilities. A helpful methodology is to ask yourself the following questions and use your knowledge of your environment and the specifics of this chapter to formulate the answers:

- ◆ What are your assets?
- ◆ What threats apply?
- ◆ What are your vulnerabilities?
- ◆ What are your resulting exposures and risk?
- ◆ How much risk can you tolerate?
- ◆ What can you afford to mitigate these risks to reduce the residual risk to a figure within your tolerance range?

APPLY YOUR KNOWLEDGE**Exercises****10.1 The Airports Council International Exercise**

The purpose of this exercise is to practice some of the concepts you learned in this chapter. Note the following:

Airports Council International shows (ACI, www.airports.org/traffic/passengers/html) that the 30 busiest airports reported the following preliminary data for passenger traffic in 2001:

<i>Rank</i>	<i>Airport</i>	<i>Number of passengers</i>
1	ATLANTA, GA (ATL)	75,849,375
2	CHICAGO, IL (ORD)	66,805,339
3	LOS ANGELES, CA (LAX)	61,024,541
4	LONDON, GB (LHR)	60,743,154
5	TOKYO, JP (HND)	58,692,688
6	DALLAS/FT WORTH AIRPORT, TX (DFW)	55,150,689
7	FRANKFURT, DE (FRA)	48,559,980
8	PARIS, FR (CDG)	47,996,223
9	AMSTERDAM, NL (AMS)	39,538,483
10	DENVER, CO (DEN)	36,086,751
11	PHOENIX, AZ (PHX)	35,481,950
12	LAS VEGAS, NV (LAS)	35,195,675
13	MINNEAPOLIS/ST PAUL, MN (MSP)	35,170,528
14	HOUSTON, TX (IAH)	34,794,868
15	SAN FRANCISCO, CA (SFO)	34,626,668
16	MADRID, ES (MAD)	33,984,413
17	HONG KONG, CN (HKG)	32,553,000
18	DETROIT, MI (DTW)	32,294,121
19	MIAMI, FL (MIA)	31,668,450
20	LONDON, GB (LGW)	31,182,361
21	BANGKOK, TH (BKK)	30,623,764
22	NEWARK, NJ (est)(EWR)	30,500,000
23	NEW YORK, NY (est)(JFK)	29,400,000
24	ORLANDO, FL (MCO)	28,166,612

continues

APPLY YOUR KNOWLEDGE

continued

<i>Rank</i>	<i>Airport</i>	<i>Number of passengers</i>
25	SINGAPORE, SG (SIN)	28,093,759
26	TORONTO, ON, CA (YYZ)	28,042,692
27	SEATTLE/TACOMA, WA (SEA)	27,036,074
28	ST LOUIS, MO (STL)	26,719,022
29	ROME, IT (FCO)	25,563,927
30	TOKYO, JP (NRT)	25,379,370
		1,166,924,477

When answering the following, concentrate on major issues. Do not try to incorporate all possible variables such as extra personnel to cover sick leave and vacation time. The rounding of calculations is appropriate (for example, use 31,000,000 rather than 31,182,361). Show your calculations and rounding.

Estimated Time: 30 minutes

1. Assume that each passenger checks one bag and that the peak passenger load in a day is five times the average load (Thanksgiving and Christmas, for example). Further assume that baggage screening machines that search for explosives and other contraband can scan 1,000 bags per hour per machine, and that these machines have a mean time between failures (MTBF) of 8,568 hours of continuous operation (not quite 1 year) and are out of service for repair for 1 week. How many baggage-screening machines are needed for the Atlanta airport to ensure that all passengers can leave on the same day they enter?
2. Each machine costs \$1.4 million; what is the capital cost Atlanta can expect for screening machines?

3. The best known device to screen baggage for explosives is a trained dog and handler. Using the same assumptions as in question 1, and further assuming that a dog can work for 2 hours at a time (during which 4,000 bags can be checked) and then needs a break of 2 hours, how many dog teams will the Atlanta airport need to ensure that all passengers can leave on the same day they enter? Assume each team works a normal 8-hour shift.
4. Discuss the use of biometrics for passenger identification. Include a discussion of error rates and mechanisms to handle errors.
5. Evaluate the answers to this exercise:

Answer to question 1: 76 million bags/year (rounded) times 5 for peak load yields an average of 208,219 bags/day, 1.041 million bags on peak days, and an average of 43,379 bags/hour. Thus, 44 machines are needed. Each machine will lose 1 week per year (44 weeks total), so a minimum of one machine (44/52) is needed to cover expected failures. The total number of baggage-screening machines is therefore 45.

APPLY YOUR KNOWLEDGE

This neglects that enough spare machines must be available at the right times in case several machines die simultaneously or things grind to a halt (probably at the peak demand time), and this makes no allowance for longer (or shorter) repair times. It also assumes an even distribution of machine use by time-of-day. (For comparison, fewer than 200 such machines were operational in the U.S. in late 2001. It would require some 700 such machines just to do 100% check-in baggage screening at only the 30 most-active airports. And there have been reports of much higher failure rates than assumed here, up to 21% down time and weeks to repair.)

Answer to question 2: \$63 million. This does not include spare part inventories, alterations to facilities to support baggage flow through those 45 machines (and the machines, which weigh around 9 tons each), time and cost to closely examine “potential problem” baggage flagged, and so on. It also does not include costs for operators and response staff, maintenance for the baggage conveyor systems, and so on. Credit yourself correct for the multiplication by a number other than 45, if it’s your output from question 1.

Answer to question 3: 189 dog/handler teams. (Dogs are twice as fast as machines but can work only 50% of a shift, so you need the same 45 teams 24/7/365. Three shifts triples that. Multiply by 7/5 to allow for 5-day work weeks.) This ignores sick and vacation time for both dogs and handlers and assumes dogs can and will do this sort of thing every day indefinitely. (Dogs won’t, and allowances for sick time and holiday time usually are around 10%–15%.) Also, fewer teams can be used if there is more than one dog per handler.

Answer to question 4: This is an open-ended essay. Note the need for the following:

- Very accurate biometric sensors
- Staff to conduct in-depth examination of passengers flagged as “suspicious”
- Time taken to recognize passengers and search a database
- Delay times in line-ups and such
- Costs
- The need for a proper risk assessment to guide selection

Bonus points for mentioning the relationship between false positive and false negative and for mentioning privacy problems. There is enough information in the chapter and in the previous table to suggest that biometric systems for passenger identification probably are not feasible; however, this should be guided by a risk assessment.

Review Questions

1. What are the three principles of physical security?
2. Name the four classes of physical assets this chapter uses.
3. List the four main types of vulnerability.
4. Vulnerabilities are further broken into _____ and _____.
5. List four general methods for controlling theft.
6. Describe a simple way to control theft of computers.

APPLY YOUR KNOWLEDGE

7. Some kinds of computer components, such as memory chips, are small, portable, and worth more than their weight in gold. How can you control theft of such things?
8. Assuming you can choose a location, what is a good way to minimize vulnerability to crime, riots, and demonstrations?
9. List at least two concerns that impact the decision to use biometrics in access control.
10. The three most common problems related to power supplies for computers are _____, _____, and _____.
11. List three types of exposure to water-related problems.
12. Clarify the difference between fire protection and fire prevention.
13. List four desirable characteristics of media storage area.
14. What is remanance, and what is the relationship of remanance to erasing media?
15. What is probably the biggest problem with installing sophisticated sensors in a perimeter detection system?
16. Describe how to perform a puff test.
17. Various types of see-through devices can display images of the contents of sealed parcels, baggage, and so forth. However, security testing at airports that use such see-through devices consistently has demonstrated poor results, with up to 50% of contraband items missed (and sometimes much higher). This might be because of an inherent problem with such devices. What is this inherent problem?

Exam Questions

1. Which of the following is probably the most common physical security issue affecting a workplace?
 - A. Theft
 - B. Destruction of company property due to floods
 - C. Terrorism
 - D. Accidental loss
2. What is a mantrap?
 - A. A device that can be deployed on the grounds of the facility and used to catch an intruder
 - B. An entrance that permits only one person at a time to pass, and that usually can be locked to trap an intruder
 - C. A special intrusion detection device that recognizes when an unauthorized individual is in the data center
 - D. In a honeypot, the part that traps the intruder and keeps him from accessing other areas of the network
3. Which two groups of people often are not considered in access control planning?
 - A. Secretaries and salespeople
 - B. Janitors and salespeople
 - C. Janitors and emergency response personnel
 - D. Contractors and emergency response personnel

APPLY YOUR KNOWLEDGE

4. What single provision covers most power supply problems, and some contingency issues as well?
 - A. UPS
 - B. Generator
 - C. Using laptops
 - D. Degausser
5. What is the single most important thing to do after detecting a fire?
 - A. Evacuate people.
 - B. Call the insurance company.
 - C. Call the fire department.
 - D. Gather critical documents.
4. Vulnerabilities are further broken into deliberate and accidental. See the “Vulnerabilities” section for more information.
5. Four ways to control theft are as follows: hire and authorize trustworthy people; make honesty part of the corporate culture; motivate people well; and minimize easy targets. See the “Vulnerabilities” section for more information.
6. There are many simple ways to control theft of computers. One way is to use a cable to attach the computer to something hard to move. Another simple way is to provide good lighting and visibility. See the “Vulnerabilities” section for more information.
7. One easy way is to lock computer cases. See the “Vulnerabilities” section for more information.
8. You can minimize vulnerability to crime, riot, and demonstrations by locating your facility near police and fire protection facilities. Also, locate it away from obvious targets (this is becoming more difficult, as obvious targets change with political shifts). See the “Selecting, Designing, Constructing, and Maintaining a Secure Site” section for more information.

Answers to Review Questions

1. The three principles of physical security are as follows: Identify the assets you need to protect; assess vulnerabilities and threats; and select countermeasures to contain the expected losses within an acceptable threshold of risk. See the “Classifying Assets to Simplify Physical Security Discussions” section for more information.
2. The four major classes of assets are facility, support, physical components, and supplies and materials. See the “Classifying Assets to Simplify Physical Security Discussions” section for more information.
3. The four main types of vulnerability are destruction; disclosure, removal, and interruption. See the “Vulnerabilities” section for more information.
4. The three most common problems related to power supplies for computers are brownouts, spikes and surges, and static. See the “Power” section for more information.

APPLY YOUR KNOWLEDGE

11. Water-related problems include flood, leaky basements, leaky roofs or drain pipes, snow loading, hurricanes, sprinkler systems, and air conditioning. See the “Water Exposure Problems” section for more information.
12. Fire protection includes detection and minimizing harm after a fire starts; fire prevention relates to avoiding the occurrence of fire in the first place. See the “Fire Prevention and Protection” section for more information.
13. Four desirable characteristics of media storage area are restricted, controlled, locked, and provided with fire prevention and protection. See the “Tape and Media Library Retention Policies” section for more information.
14. Remanance is the property of materials to retain an impression of magnetic fields after the field is removed. Its relevance to data erasure is that magnetic materials might retain a record of the data written even after normal degaussing and erasure or overwriting. Technology that allows recovery from rocks of a magnetic record of the Earth’s field reversals over hundreds of thousands of years might also allow recovery of data from erased, degaussed, or overwritten magnetic media. See the “Waste Disposal” section for more information.
15. Installing such equipment without appropriate threat and risk assessment, real-time monitoring capability, or well-defined response procedures is probably the biggest problem with installing sophisticated sensors in perimeter detection systems. See the “Physical Intrusion Detection” section for more information.
16. Open the outlets that provide a supply of fire suppression gas, cover the openings with light-weight covers, and then blow air into the system. If everything is clear, the covers over the outlets should lift or otherwise signify the free flow of air. See the “Fire Prevention and Protection” section for more information.
17. Interpreting images is difficult because many items look very different in different orientations. For example, a bottle looks like a rectangle from one angle, but like a circle from an end view. (Radiation issues can affect some devices but not all, and not even all x-ray-based devices. Personnel training problems are also a problem but are not limited to see-through scanners.) See the “Physical Intrusion Detection” section for more information.

Answers to Exam Questions

1. **A.** Probably the most common physical security issue affecting a workplace is theft. See the “Vulnerabilities” section for more information.
2. **B.** A mantrap is an entrance that permits only one person at a time to pass, and it usually can be locked to trap an intruder. See the “Selecting, Designing, Constructing, and Maintaining a Secure Site” section for more information.
3. **C.** Two groups of people who often are not considered in access control planning are janitors and emergency response personnel. See the “Active Physical Access Controls” section for more information.

APPLY YOUR KNOWLEDGE

4. **B.** A UPS filters out surges and grounds static and can completely isolate the computer from line power. It also covers the contingency of loss of power. A generator replaces power but does not deal with the surges and spikes issue.

A degausser is a demagnetizer. Laptops cannot be replacements for all computing systems and even though they include their own batteries, they are subject to damage from spikes. See the “Minimizing Power Problems” section for more information.

5. **A.** Evacuate the personnel first. See the “Fire Prevention and Protection” section for more information.

APPLY YOUR KNOWLEDGE**Suggested Readings and Resources**

1. Bruschweiler, Wallace S. Sr., "Computers as Targets of Transnational Terrorism." In *Computer Security*, edited by J. B. Grimson and H. J. Kugler. Elsevier Science Publishers, 1985.
2. "Case Histories in Computer Security." *Computer Security*. No. 53, July/August 1983.
3. *Disaster Planning for Government of Alberta Records*. Records Management Branch, Alberta Public Works Supply and Services. 10442 - 169 Street, Edmonton, Alberta T5P 3X6, 1987.
4. "EDP Threat Assessments: Concepts and Planning Guide." *RCMP Security Information Publications # 2*. January 1982.
5. Emergency Preparedness Canada. *Guide to the Preservation of Essential Records*. EPC 12/87, December 1987.
6. Fites, Philip. E., P. Martin, J. Kratz, and Alan F. Brebner. *Control and Security of Computer Information Systems*. New York: W. H. Freeman/Computer Science Press, 1989.
7. Gallegos, Frederick, Dana R. (Rick) Richardson, and A. Faye Borthick. *Audit and Control of Information Systems*. Chicago: Southwestern Publishing Company, 1987.
8. "Good Security Practices for Personal Computers." *IBM Data Security Support Programs, First Edition*. 1984.
9. Jacobson, Robert V., et al. "Guidelines for Automatic Data Processing Physical Security and Risk Management." *Federal Information Processing Standards Publication 31*. National Bureau of Standards, 1974.
10. Lobel, J. *Foiling the System Breakers: Computer Security and Access Control*. New York: McGraw-Hill, 1986.
11. Parker, Donn B. *Computer Security Management*. Reston, Virginia: Reston Publishing Company, Inc., 1981.
12. Parker, Donn B. *Fighting Computer Crime*. New York: Charles Scribner and Sons, 1983.
13. *Personal Computer Security Considerations* (NCSC-WA-002-85). National Computer Security Center. Ft. George G. Meade, Maryland: December 1985.
14. "Small Computer Systems Security," and "Small Systems Questionnaire." In *EDP Security Bulletin, RCMP "T" Directorate, Vol. 12 No. 1*. July 1987. (The questionnaire is not copyrighted and may be reproduced for use; it is also in French and English.)
15. "Target Hardening." *RCMP Security Information Publications # 3*. September 1983.



P A R T



II





FINAL REVIEW

Fast Facts

Study and Exam Preparation Tips

Practice Exam





The Fast Facts listed in this chapter are designed as a refresher of key points and topics, knowledge of which is required to be successful on the CISSP Certification exam. By using these summaries of key points, you can spend an hour prior to your exam to refresh your understanding of key topics and ensure that you have a solid understanding of the information required for you to succeed in each domain of the exam.

The chapter is organized by domains. It is designed as a quick study aid you can use just before taking the exam. You should be able to review the Fast Facts in less than an hour. It cannot serve as a substitute for knowing the material supplied in these chapters; however, its key points should refresh your memory on critical topics. In addition to the information located in this chapter, remember to review the Glossary terms because they are intentionally not covered here and are important to your understanding of the material.

The following list shows the domains discussed in this book, which is how we determined a set of objectives for each domain:

- ▶ Domain 1, "Access Control"
- ▶ Domain 2, "Network Security and Telecommunications"
- ▶ Domain 3, "Security Management and Practices"
- ▶ Domain 4, "Applications and Systems Development Security"
- ▶ Domain 5, "Cryptography"
- ▶ Domain 6, "Security Architecture and Models"
- ▶ Domain 7, "Operations Security"
- ▶ Domain 8, "Business Continuity Planning and Disaster Recovery Planning"
- ▶ Domain 9, "Law, Investigation, and Ethics"
- ▶ Domain 10, "Physical Security"



Fast Facts

DOMAIN 1, “ACCESS CONTROL”

Accountability

Except in certain extreme circumstances, shared accounts must be avoided. This policy must be clearly reflected in the security policy and strictly enforced.

Access Controls

The typical access control types used are

- ◆ **Discretionary access control (DAC)**—Essentially based on human decisions about whether someone (or a service, an application, and so on) should be allowed access to a particular resource, such as a file or directory.
- ◆ **Mandatory access control (MAC)**—Applies a higher level of access control in which the computer system strictly controls who can access what resources. Because MAC is based on using classification levels, it is more popular in government-type environments. You should be aware that an issue common in MAC environments is that users could have multiple accounts associated with different levels of access. This is a limitation because a user might log on with the highest level to do all his work. A solution is the use of a MAC, which does not allow access to lower-level areas by higher-level access user accounts.
- ◆ **Lattice-based access control**—A form of MAC for strictly implementing access controls across an organization. A set of security classes and a set of flow operations are defined. Flow operations determine when information can flow from one class to the other.

- ◆ **Rule-based access control**—Involves setting up parameters around which an individual can access a system. This type of access control system does not scale well.
- ◆ **Role-based access control**—You develop roles or positions across your company and assign access to the role based on the job functions of that position. This is the most widely used form of access controls.
- ◆ **Access controls lists (ACLs)**—Used to create a list of rules, perhaps based on IP addresses or some piece of information that can easily be discernible in the packets that go across the network. For each rule, you specify whether you will allow or deny traffic. ACLs are often associated with routers and applied to limit the amount of traffic that can go to a given network resource. ACLs are also used in file systems to assign access such as read, write, execute, and delete.

Access Control Administration

Administering access control includes the following:

- ◆ Assigning account IDs and passwords for users
- ◆ Managing accounts by assigning permissions to accounts
- ◆ Assigning and maintaining an account policy that might include rules to control passwords, logon times, and so on

Access Control Models

The Bell-LaPadula (BLP) model deals with the flow of information from a confidential standpoint. BLP is composed of two rules:

- ◆ Simple security deals with reading information or files.
- ◆ The star property deals with writing information or creating new files.

The Biba model is similar to BLP except for the fact that, instead of dealing with confidentiality, it deals with integrity. BLP has the following two rules:

- ◆ Simple security deals with reading.
- ◆ The star property deals with writing.

The big difference, which seems confusing at first, is that both of the rules are the opposite of the BLP model.

The Liptner model applies lattices and the principals of integrity and confidentiality to non-military examples. Essentially, Liptner changed the labels from terms such as *confidential* and *secret* to *system programmers*, *production code*, and so on.

The non-interference models deal with examining the input to and output from a system and seeing whether you can infer any information that you should not have access to.

Identification and Authentication Techniques

Identification is a statement of who you are, such as a user ID or logon name. Authentication is proving you are who you say you are. Several techniques are used by systems to provide authentication:

- ◆ Passwords
- ◆ One-time passwords
- ◆ Challenge response
- ◆ Biometrics

- ◆ Tickets
- ◆ Single sign-on (SSO)

Three things represent techniques that can be used for authentication:

- ◆ **Something you know**—Passwords
- ◆ **Something you have**—One-time passwords
- ◆ **Something you are**—Biometrics

Remote Authentication Access Control

RADIUS and TACACS+ are typically used interchangeably for remote access controls.

Centralized Versus Decentralized Access Control

With *centralized control*, a single authority or system is responsible for access control. The biggest problem with this is that a single point of failure exists that could also become a bottleneck for an organization. With decentralized control, each individual or department is responsible for its own access control.

Methods of Attack

Types of attacks include

- ◆ **Brute-force**—With a *brute-force attack*, an intruder tries all possible combinations until she guesses the right one. Brute-force attacks are most popular when cracking passwords.
- ◆ **Denial-of-service**—These involve preventing others from gaining access.

578 FAST FACTS

- ◆ **Spoofing**—An attacker acquires the one-time password device (or other appropriate access control process) for a given user and acts like that user (or *spoofs* that user). The system then gives her access because the system thinks she is a legitimate user and does not know that she is really an attacker.
- ◆ **Sniffing**—The process of capturing the packets traveling across the wire and either reading plaintext passwords or capturing credentials and cracking them.

Monitoring

The field of study dealing with monitoring networks and hosts and looking for attacks is known as *intrusion detection*. The critical thing to remember with intrusion detection is that you are passively monitoring a network or hosts looking for signs of an attack. The emphasis is on detection, not prevention.

Signature or pattern matching IDS maintains a database of known attack signatures. When it looks at traffic or at log files, it tries to find a match for each of these signatures. If it finds a match, it sends an alert that the system is being attacked.

The concept behind anomaly detection is to determine what is normal traffic for a company, and anything that falls outside that norm is deemed an attack and is dropped.

Penetration Testing and System Assessment

Penetration testing is sometimes contrasted or compared with *security assessments*. The main difference between the two has to do with the scope and amount of initial information one is given. Typically, with a penetration test (or *pen test*), the goal is to see how much you can find out about the company, including possible ways you can break in.

Security assessments usually include a penetration test but are much more thorough. You are typically given access to all the key systems within a company to evaluate the current level of security. With security assessments, you are not trying to prove that you can get in; you are trying to paint a picture of the current threats that exist to the organization and what needs to be done to protect against them.

DOMAIN 2, “NETWORK SECURITY AND TELECOMMUNICATIONS”

ISO/OSI Seven-Layer Model

The ISO/OSI seven-layer model defines the fundamental aspects of how all network communication occurs. The OSI model exists to enable the user to understand the totality of a very complex system of communications by breaking the overall transmission of data into seven easier-to-define layers:

- ◆ **Application layer**—Primarily responsible for interfacing with the user. This is the application interface the user experiences. (POP3, NNTP)
- ◆ **Presentation layer**—Primarily responsible for translating the data from something the user expects to something the network expects. (WAV, MIDI, JPEG, SMB)
- ◆ **Session layer**—Primarily responsible for dialog control between systems and applications. (NSF, RPC)
- ◆ **Transport layer**—Primarily responsible for handling end-to-end data transport services. (TCP, UDP)

- ◆ **Network layer**—Primarily responsible for logical addressing. (IP, IPX)
- ◆ **Data Link layer**—Primarily responsible for physical addressing. (IEEE 802.2, 802.3, switches, bridges)
- ◆ **Physical layer**—Primarily responsible for physical delivery and specifications.

Network Cabling

This section looks at cable specifications for coax, UTP, fiber, and wireless.

Coax

The following cable specifications exist for coax cable:

- ◆ **RG-58 /U**—Solid copper core (0.66mm or 0.695mm), 53.5 ohms.
- ◆ **RG-58 A/U**—Stranded copper core (0.66mm or 0.78mm), 50 ohms.
- ◆ **RG-58 C/U**—Military version of RG58 A/U (0.66mm), 50 ohms
- ◆ **RG-59**—Broadband transmissions (for example, cable TV)
- ◆ **RG-6**—Higher frequency broadband transmissions; a larger diameter than RG-59
- ◆ **RG-62**—ArcNet
- ◆ **RG-8**—Thicknet, 50 ohms

Coax networks are less commonly used than 10BASE-T networks because coax has a single point of failure for the entire segment and is more difficult to troubleshoot.

The maximum number of nodes per segment (between repeaters) on a 10BASE-2 segment is 30. The maximum length of a segment is 185 meters. You can actually determine the maximum cable length by the name 10BASE-2. 10 stands for 10Mbps; BASE stands for baseband; and 2 stands for 200 meters (okay, so it is a little short).

10BASE-5 supports a maximum of 1,024 hosts per segment. The maximum segment length for 10BASE-5 is 500m.

10BASE-2 and 10BASE-5 adhere to the 5-4-3 rule. This simply means that you can have a maximum of five segments connected via four repeaters, but only three segments can have hosts on them. The two segments that cannot support hosts are called interrepeater links (IRL).

A time domain reflectometer (TDR) can be used on one end of the cable to give an approximate distance within a few feet or so to the break in the wire.

UTP Cabling

The most common type of cabling for Ethernet LANs is UTP. UTP cable comes in 10BASE-T and 100BASE-TX media types. The 10 and 100 refer to the speed the network runs at—either 10Mbps or 100Mbps. The cabling specification for this topology is known as Category 3, 4, 5, 5E, 6, and 7. The category of cabling indicates the quality of the signal carrying, as well as the number of wires used and number of twists in the wires.

The following are category and speed ratings for UTP cables:

Category 3	Rated for voice and data up to 10Mbps/16MHz
Category 4	Rated for voice and data up to 16Mbps/20MHz

580 FAST FACTS

Category 5	Rated for voice and data up to 100Mbps/100MHz (most widely used at present)
Category 5e	Rated for voice and data up to 1000Mbps/100MHz
Category 6	Rated for voice and data up to 1000Mbps/250MHz
Category 7 (proposed draft)	Rated for voice and data up to 10000Mbps/600MHz

10BASE-T *unshielded twisted pair (UTP)* cabling has no shielding, and the four pairs of conductors twist around each other inside the cable jacket. Because there is no shielding, UTP is very susceptible to electromagnetic interference (EMI), such as the EMI given off by fluorescent lights. UTP also enables a malicious user to easily capture the data being transmitted without ever needing to tap into the cable.

Fiber

Fiber-optic cable is predominately used for backbone and device interconnectivity, as opposed to end user connectivity. The individual fiber strands are then typically bundled in pairs or multiple pairs because each fiber can send a signal in only a single direction.

Wireless

There are a few rather substantial drawbacks to wireless at this time. The first is the lack of standardization; the other problem with wireless is one of security. Just as anyone can tune his radio to receive certain radio stations, people can connect to a wireless network by simply running the appropriate equipment and being within a certain range.

Network Topologies

Virtually all networks use one of the following topologies:

- ◆ **Linear Bus**—All the systems are connected in a row to a single cable in a daisy-chain fashion (coax).
- ◆ **Star**—Unlike coax, the topology method in a 10BASE-T network is a star because all the devices must have a segment of wire connecting them to an active hub or switch before being capable of communicating with other devices on the LAN.
- ◆ **Ring**—Designed using a loop of cable to interconnect the devices. The signal is transmitted in a single direction around the loop, with each device retransmitting the signal as it receives it.
- ◆ **Tree**—Based in part on the bus and the star topology. In the tree topology, devices are interconnected to each other via bus connections; however, multiple nodes are supported on each potential branch.
- ◆ **Mesh**—Every node on a network is connected to every other node.

LAN and WAN Technologies

Transmission techniques consist of the following:

- ◆ **Unicast**—The packet is addressed to a specific destination host, both physically and logically.
- ◆ **Broadcast**—The packet is destined to all hosts on a subnet or network. At the Data Link layer, the address used is FFFFFFFF (All Fs) in hexadecimal. At the network layer, the address used is the network broadcast identifier—or the all networks broadcast address of 255.255.255.255.

- ◆ **Multicast**—The packet is addressed to multiple hosts via the use of group membership addresses.

Ethernet

Ethernet is the single most predominant technology in use today, with speeds ranging from 10Mbps to 10Gbps. Ethernet uses CSMA/CD, which helps the devices on the network share the bandwidth while ensuring that two devices cannot use the bandwidth at the same time.

Ring Topology

The most predominant method of transmitting data on a ring topology is through the use of something called *token passing*. The token is simply a packet to which data is appended for transmission. As a result, if a system wants to transmit, it must have the token so that it can append the data to the token and transmit it.

Network Devices

Following are network devices:

- ◆ **Hubs and repeaters (Physical layer)**—The primary functions of a hub (repeater) are to receive a signal, amplify the signal, and repeat the signal out all ports.
- ◆ **Switches and bridges (Data Link layer)**—Switches read at least part of the data and attempt to determine to which port the destination host is connected. If the switch can determine the destination port, it sends the signal only on the destination port. A Layer 3 switch is simply a hybrid device that combines Layer 2 and Layer 3 functionality, allowing the switch to forward frames when possible and route packets when needed. Bridges are similar to switches.

One major difference, however, is that a bridge can run only one instance of spanning tree, whereas switches can have multiple instances. Spanning tree is a protocol, defined in the IEEE 802.1d standard, that is responsible for preventing loops from occurring on a bridged/switched network.

- ◆ **Virtual LANs (VLANs)**—The creation of logically segmented networks within a single switch or within a single switch fabric. A *switch fabric* is a group of switches that are physically connected to each other.
- ◆ **Routers (Network layer)**—These can further optimize network traffic by using the logical addressing information available from the Network layer. Routers are considered “network aware,” which means routers can differentiate between different networks.

Firewalls

Firewalls are designed to prevent unauthorized traffic from entering a network. They are typically deployed as a perimeter security mechanism to screen Internet traffic attempting to enter the network. The following are the types of firewalls:

- ◆ **Packet filtering firewalls**—Function by comparing received traffic against a ruleset that defines what traffic is permitted and what traffic is denied
- ◆ **Application filtering firewalls**—Function by reading the entire packet up to the Application layer before making a filtering decision
- ◆ **Stateful inspection firewalls**—Track the network connection state and then use it in determining what traffic should be allowed to pass back through the firewall

Gateways and Proxies

In its most basic definition, a gateway provides access to a network or service. Proxies are used as intermediary devices between a client and a server, providing the client transparent access to the resources on the server without allowing the client to access those resources directly.

Connection Speeds and Types

The more common connection speeds and types are

- ◆ **Digital Signal Level 0 (DS-0)**—Defines the framing specification used to transmit data on a single 64Kbps channel over a T1 line.
- ◆ **Digital Signal Level 1 (DS-1)**—Defines the framing specification for transmitting data at 1.544MBps over a T1 or 2.048Mbps on an E1 line.
- ◆ **Digital Signal Level 3 (DS-3)**—Defines the framing specification for transmitting data at 44.736Mbps on a T3 line.
- ◆ **T1**—A T1 carries 24 PCM (pulse code modulations) signals, sometimes called channels, using TDM (time division multiplexing) to achieve a transmission speed of 1.544MBps over a dedicated connection.
- ◆ **T3**—A T3 carries 672 PCM signals, sometimes called channels, using TDM to achieve a transmission speed of 44.736Mbps over a dedicated connection.
- ◆ **E1**—Similar to a T1, E1s are used primarily in Europe and carry data at 2.048Mbps.
- ◆ **E3**—Similar to an E1, E3s are used primarily in Europe and carry data at 34.368Mbps.
- ◆ **OC-x (Optical Carrier X)**—The various optical carriers are a subset of the SONET (Synchronous Optical Network) specification for transmitting digital signals over fiber-optic cable. The base OC rate of OC-1 is 51.84Mbps. The numeric value of the OC rate is multiplied by the base rate to get the speed. OC-3 transmits at 155.52Mbps, OC-12 is 622.08Mbps, OC-24 is 1.244Gbps, OC-48 is 2.488Gbps, OC-192 is ~10Gbps, OC-256 is 13.271Gbps, and OC-768 is ~40Gbps.

Connections

Three types of device connections are

- ◆ **Circuit switched**—When two devices need to communicate with each other, the data network they are using dynamically brings up the circuits (or connections) the two devices require to exchange data.
- ◆ **Packet switched connections**—Use a synchronous serial method of communications. Where packet switching differs is that the packet switched network is often shared by multiple systems.
- ◆ **Cell switched networks**—Very similar to packet switched networks with one important difference: Cell switched networks are Asynchronous Transfer Mode-based networks. Asynchronous Transfer Mode (ATM) is a networking standard that uses fixed length 53-byte cells in the transmission of multiple services, such as voice, video, and data.

WAN Services

Multiple services can be used for communication on a wide area network (WAN):

- ◆ **Serial Line Internet Protocol (SLIP) and Point-to-Point Protocol (PPP)**—Primarily used for providing data-link connectivity over asynchronous (dial-up) and synchronous serial (ISDN or dedicated serial lines such as T1) connections.
- ◆ **High-Level Data Link Control (HDLC)**—An ISO-based standard for delivering data over synchronous serial lines.
- ◆ **X.25**—A WAN connection technique that functions at the Physical and Data Link layers of the OSI model. X.25 uses virtual circuits for establishing the communications channel between hosts.
- ◆ **Frame relay**—Based on X.25; however, it is considered a faster technology because it leaves error correcting functionality to higher layers.
- ◆ **Synchronous Data-Link Control (SDLC)**—A bit-oriented connection protocol designed by IBM for use in mainframe connectivity.
- ◆ **Integrated Services Data Network (ISDN)**—Developed as a standard for transmitting digital signals over standard telephone wires.
- ◆ **x Digital Subscriber Line (xDSL)**—A relatively new technology that supports the broadband transmission of data at high speeds, currently up to about 53Mbps, over the existing telephone network.
- ◆ **Switched Multimegabit Data Service (SMDS)**—A high-speed packet switching technology for use over public networks. It is provided for companies that need to send and receive large amounts of data on a bursty basis.

- ◆ **High-Speed Serial Interface (HSSI)**—Sometimes called “hissy,” it provides for an extremely fast point-to-point connection between devices; however, the distance limitation is no more than 50 feet.

WAN Devices

Devices used on the WAN are

- ◆ **Routers**—Although LAN devices, routers are also used extensively on WANs to provide routing between subnets.
- ◆ **WAN switches**—Used to connect private data over public circuits.
- ◆ **Multiplexors**—Enable more than one signal to be transmitted simultaneously over a single circuit.
- ◆ **Access servers**—Often are used for dial-in and dial-out access to the network.
- ◆ **Modems**—Convert digital and analog signals, allowing digital data to be transmitted over analog phone lines.
- ◆ **CSU/DSU (channel service unit/data service unit)**—Digital interface devices that are used to terminate the physical connection on a DTE device to the DCE.

Remote Access

Important things to remember about remote access are

- ◆ Tunneling is the process of transmitting one protocol encapsulated within another protocol.

584 FAST FACTS

- ◆ A VPN is simply the use of a “tunnel,” or secure channel, across the Internet or other public network. The data within the tunnel is encrypted, thus providing security and integrity of the data against outside users.
- ◆ The protocols used in VPN are Point to Point Tunneling Protocol (PPTP), Internet Protocol Security (IPSec), and Layer 2 Tunneling Protocol (L2TP).
- ◆ Remote Authentication Dial-In User Server (RADIUS) is a User Datagram Protocol-based de-facto industry standard for providing remote access authentication via a client/server model.
- ◆ Similar in function to RADIUS, Terminal Access Controller Access Control Service (TACACS+) differentiates itself by separating the authentication and authorization capabilities, as well as by using TCP for connectivity. As a result, TACACS+ is generally regarded as being more reliable than RADIUS.

TCP/IP

TCP/IP is actually a suite of protocols that was developed by the Department of Defense to provide a highly reliable and fault-tolerant communications infrastructure. It uses a four-layer model:

- ◆ **Application layer**—Loosely maps to the top three layers of the OSI model and provides for the applications, services, and processes that run on a network (BOOTP, FTP, POP3).
- ◆ **Transport layer**—Sometimes referred to as the Host-to-Host layer, the Transport layer is responsible for handling the end-to-end data delivery on the network. It loosely maps to the Transport layer of OSI (TCP and UDP).
- ◆ **Internet layer**—Maps loosely to the Network layer of the OSI model and provides for logical addressing and routing of IP datagrams on the network. (IP, ICMP, ARP).
- ◆ **Network layer**—Maps loosely to the Data Link and Physical layers of the OSI model. The Network layer is primarily responsible for the physical delivery of data on the network.

Common Network Attacks and Countermeasures

Several common network attacks are

- ◆ Social engineering
- ◆ Brute-force
- ◆ Non-business use of systems
- ◆ Network sniffing, dumpster diving, and keylogging
- ◆ Denial-of-service
- ◆ Spoofing, Trojans, viruses and worms, and backdoors
- ◆ Scanning

Fault Tolerance

Several methods are available for adding fault tolerance to networks and network devices. In a given environment, many can be used.

RAID

Fault tolerance helps mitigate the threat of disk failure. There are five levels of RAID:

- ◆ **RAID 0**—Used to provide a performance increase by allowing simultaneous reads and writes through striping of data across multiple disks.

RAID 0 provides no fault tolerance. If one disk fails, the data on all disks is lost.

- ◆ **RAID 1**—Also called mirroring, it duplicates the data on one disk to another disk.
- ◆ **RAID 2**—Uses multiple disks and parity information. Parity keeps track of whether data has been lost or overwritten by use of a parity bit.
- ◆ **RAID 3–4**—RAID 3 performs byte-level striping, and RAID 4 performs block-level striping across multiple drives. Parity information is stored on a specific parity drive.
- ◆ **RAID 5**—Stripes data and parity across all drives using interleave parity for data re-creation. Because reads and writes can be performed concurrently, RAID 5 offers a performance increase over RAID 1.

Clustering

In a data clustering scenario, the administrator configures two servers as mirrors of each other, both sharing access to a common storage system. If one of the servers fails, the services running on that server can be transferred to the backup server.

Network services clustering is used to improve system performance by distributing network requests among multiple servers that typically have the same functionality.

Backup

Backup methods include

- ◆ **Full backup**—A full backup saves every file, every time.
- ◆ **Incremental backup**—Only backs up the data that has been changed or added recently.

- ◆ **Differential backup**—A differential backup backs up files that have changed since the last full backup.

DOMAIN 3, “SECURITY MANAGEMENT AND PRACTICES”

CIA Triad

Following describes the CIA Triad (confidentiality, integrity, availability):

- ◆ **Confidentiality**—Determines the secrecy of the information asset. The level of confidentiality determines the level of availability that is controlled through various access control mechanisms.
- ◆ **Integrity**—Provides the assurance that the data is accurate and reliable.
- ◆ **Availability**—The ability of the users to access an information asset.

Privacy

Privacy relates to all elements of the CIA Triad. It considers which information can be shared with others (confidentiality), how that information can be accessed safely (integrity), and how it can be accessed (availability).

Identification and Authentication

Identification provides the resource with some type of identifier of who is trying to gain access.

586 FAST FACTS

Authentication is proving you are who you say you are. The following are some things used to do so:

- ◆ What the entities know, such as a personal identification number (PIN) or password
- ◆ What the entities have, such as an access card, a smart card, or a token generator
- ◆ Who or what the entity is, which is usually identified through biometrics

Auditing

Systems and security administrators can use the audit records to

- ◆ Produce usage reports
- ◆ Detect intrusions or attacks
- ◆ Keep a record of system activity for performance tuning
- ◆ Create evidence for disciplinary actions or law enforcement

Accountability

Accountability is created by logging the events with the information from the authenticated user.

Risk Management and Analysis

Risk management is the process of assessing risk and applying mechanisms to reduce, mitigate, or manage risks to the information assets.

Risk analysis identifies a risk, quantifies the impact, and assesses a cost for mitigating the risk.

The process of quantitative risk analysis consists of several steps, including identifying the assets, assigning value to them, identifying threats and risks, and determining how much money would be lost if the threat became reality. Potential monetary loss can be calculated using the following formulas:

- ◆ Single-loss expectancy (SLE) is the amount of the potential loss for a specific threat.
- ◆ Estimate annual frequency of occurrence or exposure factor (EF).
- ◆ Risk analysis is based on the loss over the course of a year. The *annualized rate of occurrence (ARO)* is the ratio of the estimated possibility that the threat will take place in a 1-year time frame. The ARO can be expressed as 0.0 (if the threat will never occur) through 1.0 (if the threat will always occur).
- ◆ Determine the annualized loss expectancy (ALE). Do this with the following steps:
 1. The SLE is calculated by multiplying the value of the asset by the EF:

$$\text{SLE} = \text{asset value} \times \text{EF}$$
 2. The ALE is calculated by multiplying the SLE by the ARO:

$$\text{ALE} = \text{SLE} \times \text{ARO}$$

Qualitative Risk Analysis

A *qualitative risk analysis* is a more subjective analysis that ranks threats, countermeasures, and their effectiveness on a scoring system rather than by assigning dollar values.

Cost-Effectiveness of a Countermeasure

Determining the most cost-effective countermeasure is called a *cost/benefit analysis*.

A cost/benefit analysis looks at the ALE, the annual cost of the safeguard, and the ALE after the countermeasure is installed to determine whether the costs show a benefit for the organization. The calculation can be written as follows:

$$\text{Value of Countermeasure} = \text{ALE (without countermeasure)} - \text{Cost (safeguard)} - \text{ALE (with countermeasure)}$$

Responses to Risk Analysis

After a risk analysis has been completed, an organization must choose its response. Its choices are

- ◆ **Do nothing**—If you do this, you must accept the risk and the potential loss if the threat occurs.
- ◆ **Reduce the risk**—You do this by implementing a countermeasure and accepting the residual risk.
- ◆ **Transfer the risk**—You do this by purchasing insurance against the damage.

Residual risk is the value of the risk after implementing the countermeasure.

Policies

Information security policies are high-level plans that describe the goals of the procedures. Policies are not guidelines or standards, nor are they procedures or controls. Policies describe security in general terms, not specifics.

They provide the blueprints for an overall security program just as a specification defines your next product. The following are important things to remember:

- ◆ Information security policies are the blueprints, or specifications, for a security program.
- ◆ Standards are the mandatory mechanisms to implement the information security policies.
- ◆ Baselines are the minimum levels of security that will meet policy requirements.
- ◆ Guidelines are recommendations as to how to meet policy requirements.
- ◆ Procedures describe exactly how to implement countermeasures.

Protection Mechanisms

Protection mechanisms are used to enforce layers of trust between security levels of a system. They are as follows:

- ◆ **Layering**—Used to separate resources of a system into security zones.
- ◆ **Abstraction**—The collection of data and methods managed as objects.
- ◆ **Data hiding**—Data is hidden and inaccessible from the other layers.
- ◆ **Encryption**—The conversion of data to something unreadable using a mathematical equation. Significant to that equation is a key that is used as a secret value to perform the function.

Data Classification

Table 1 describes the classifications of data.

588 FAST FACTS

TABLE 1
DATA CLASSIFICATION

<i>Classification</i>	<i>Description</i>
Sensitive	Data that is to have the most limited access and requires a high degree of integrity. This is typically data that will do the most damage to the organization should it be disclosed.
Confidential	Data that might be less restrictive within the company but might cause damage if disclosed.
Private	Private data is usually compartmental data that might not do the company damage but must be kept private for other reasons. Human resources data is one example of data that can be classified as private.
Proprietary	Proprietary data is data that is disclosed outside the company on a limited basis or contains information that could reduce the company's competitive advantage, such as the technical specifications of a new product.
Public	Public data is the least sensitive data used by the company and would cause the least harm if disclosed. This could be anything from data used for marketing to the number of employees in the company.

Government Data Classification

Table 2 describes the classifications of government data.

TABLE 2
GOVERNMENT DATA CLASSIFICATION

<i>Classification</i>	<i>Description</i>
Top Secret	Disclosure of top secret data would cause severe damage to national security.
Secret	Disclosure of secret data would cause serious damage to national security. This data is considered less sensitive than data classified as top secret.

Classification

Confidential

Description

Confidential data is usually data that is exempt from disclosure under laws such as the Freedom of Information Act but is not classified as national security data.

Sensitive But Unclassified (SBU)

SBU data is data that is not considered vital to national security, but its disclosure would do some harm. Many agencies classify data they collect from citizens as SBU. In Canada, the SBU classification is referred to as protected (A, B, C).

Unclassified

Unclassified is data that has no classification or is not sensitive.

DOMAIN 4, “APPLICATIONS AND SYSTEMS DEVELOPMENT SECURITY”**Centralized and Distributed Systems**

Systems can be centralized, distributed, or some blending of the two. Most systems fall into the following categories:

- ◆ **Centralized**—All computing takes place in one place. The old mainframe/data center approach is one example; another is the use of a mini-computer or mini-computers located in one place and held under the central control of one department. A single PC, used to support recordkeeping or other computing at a small company, can also be considered as centralized computing.

- ◆ **Centrally controlled computing**—In this scenario, computers can exist in a widely distributed fashion both within headquarters and at remote offices. They are, however, configured, maintained, and controlled by a central authority.
- ◆ **Decentralized**—Computing facilities exist throughout the company. They might or might not be linked with each other.
- ◆ **Distributed**—Computers are everywhere, and so is the process of processing. Distributed computing does not preclude centralized control.

Risks for Centralized Computer Systems

Centralized systems are vulnerable to risks specific to the type of system they are—for example, big iron (mainframe) versus mini-computer versus standalone PC.

Risks for Big Iron

Mainframe systems have their own set of risks to data, including

- ◆ Incorrect data entered in error.
- ◆ Incorrect data entered on purpose.
- ◆ Someone could enter code, which when it was run, it extracted data, modified data, destroyed data, and disrupted the systems operation.
- ◆ Unauthorized access to data either by getting past the controls (password sharing, password cracking, social engineering) or by seeing data displayed on screens in offices.
- ◆ Unauthorized use of unattended terminals where sessions are left active.

Additional Risks for Standalone PCs

PCs are also subject to the risks to data that mainframes have. In addition, they are subject to the following risks:

- ◆ Virus
- ◆ Trojan
- ◆ Logic bomb

Distributed Systems Issues

Distributed systems also can be subject to the previous risks. In addition, the following risks are present:

- ◆ Worms
- ◆ ActiveX/Java applets
- ◆ Blended malware
- ◆ Remote administration programs

Database Management Systems

Database systems have unique characteristics. Important characteristics are those specific to database management systems, database models, and database issues.

Database Management Systems

The unique characteristics that identify a database management system are

- ◆ **Data independence**—Although software is provided to assist in the management of the DBMS, the software written to provide functionality for its owners does not have to be the sole user of the data. A different program can be written to use the data.

590 FAST FACTS

- ◆ **Minimal data redundancy**—Instead of storing data in multiple places, DBMSs make data available from multiple places.
- ◆ **Data reuse**—Data gathered for one purpose can be mined for use in another.
- ◆ **Data consistency**—Data viewed or retrieved in different ways will be the same.
- ◆ **Persistence**—The state of the database and its data remains the same after code is executed.
- ◆ **Data sharing**—Many users can access the database at the same time.
- ◆ **Data recovery**—In the event of an error or a system crash, the system can recover.
- ◆ **Security controls**—A database should be capable of providing variable security controls by limiting access to those who require it.
- ◆ **Data relationships defined by primary and foreign keys**—The *primary key* of a table is the data field or column that is used as the primary index and that allows a relationship to be built with another file.
- ◆ **Data integrity consisting of semantic and referential integrity**—*Semantic integrity* is enforced by rules that specify constraints.
- ◆ **Utilities or processes to ensure efficient processing overtime**—These include *compression* (the capability to compress data and save storage space and I/O), reorganization or *defragmentation* (reclaiming of unused space), and *restructuring* (the capability to add and change records, data, access controls, disk configuration, and procedural methods).

Database Models

Not all database systems are the same. The major classifications or models of database systems are as follows:

- ◆ **Relational**—In these database models, data is stored in tables that consist of rows (like records in a regular file) and columns (like fields). Relationships are formed between tables based on a selected primary key.
- ◆ **Hierarchical**—Data is organized in a tree structure with a tree being composed of branches, or *nodes*. Think of the branches as data records, and think of the leaves of the branches as the data.
- ◆ **Network (IDMS/R)**—Data is represented in blocks or record types. Blocks include data fields, and arrows between the blocks represent a relationship between the data.
- ◆ **Object-oriented**—Combines the object data model of object-oriented programming with DBMS.
- ◆ **Distributed**—In the typical databases (object-oriented, relational, and so on), data resides on one computer. In the distributed model, data can be partitioned across multiple computers and locations.

Database Issues

Issues that can cause security issues with database systems are

- ◆ Default administrative passwords.
- ◆ Misuse, or no use, of test database.
- ◆ Lack of separation of data administration from application system development.
- ◆ Distributed databases have multiple access points.

- ◆ Distributed database processing is much harder to get right.
- ◆ Aggregation of data can expose sensitive information.
- ◆ Denial-of-service attacks.
- ◆ Improperly modifying data.
- ◆ Access to some data can provide the ability to deduce or infer data that is protected.

SANs

Benefits of SANs include

- ◆ Centralized control, including backup and management.
- ◆ Access from anywhere at anytime.
- ◆ Can improve data protection.
- ◆ Additional storage can be added with little to no disruption.
- ◆ Better physical security.
- ◆ Improved availability.
- ◆ Business flexibility.
- ◆ Can improve disaster tolerance.

To secure SANs networks, you must do the following:

- ◆ Centralize storage.
- ◆ Require encryption when IP is used.
- ◆ Authenticate users.
- ◆ Implement access controls.

Web Services Issues

Some issues specific to Web services are

- ◆ Security between vendor-specific models.
- ◆ Processing is transparent. This ensures little notice of activity by end users and administrators and can obscure security issues.

More on Attacks

The following are two attacks unique to software:

- ◆ If an instruction is executed in more than one step, it might be possible to compromise the system by attacking between the steps. Time of Check to Time of Use (TOC/TOU) is the name for a special type of race condition that can be vulnerable to this type of attack. IBM's OS 360 (an older mainframe system) performed access control over files by first reading and checking permissions; then, if the permissions were correct, the file would be read again. If the permission were incorrect, the user would be denied access. However, if the system could be interrupted before the denial was returned, the file could be read and possibly modified. More recent race conditions (conditions that exist because of timing issues within software) include problems with the `rm` command in Linux. Because of the way the command was written, it could be reissued before complete, causing a DoS for an unprivileged user and a possible removal of the entire file system if the user was a root user. This error is not present in updated versions of the OS.
- ◆ Illegitimate use of remote access software.

Malicious Code

Malicious code is any code that, either by design or as the result of being run, accomplishes any of the following:

- ◆ Modifies computer programs without the consent of the owner or operator
- ◆ Crashes programs or systems
- ◆ Steals or modifies data
- ◆ Inserts or adds code to a system that might do damage later

System Development Models

Common models used in systems development are waterfall, spiral, and rapid application development.

Waterfall Model

The classical waterfall model approach to software development has been around for a very long time. Each step from conceptual development to maintenance flows from the top down:

- ◆ **Conceptual Definition/Feasibility Study**—The need for the software to be developed is described and flushed out during an initial discovery phase.
- ◆ **Systems Analysis/Functional Requirements Determination**—Precise descriptions of exactly what is needed. This is done to a fine, granular level of detail.
- ◆ **Design/Specifications Development**—A detailed design of how the system will look. It is said that if this is done well, the pseudocode (precise descriptions of the processing with no programming language used) can easily be converted into code with little modification.

- ◆ **Design Review**—A step-by-step review of the design, measuring it against the functional specification.
- ◆ **Construction**—The program is coded according to the design.
- ◆ **Code Review or Walk-through**—Code is reviewed in excruciating detail, step by step, to ensure the program matches the design.
- ◆ **System Test Review**—All aspects of the code are tested looking for functionality, design flaws, and bugs.
- ◆ **Certification/Accreditation**—If the code must meet or is scheduled to meet some formal review for certification or accreditation, this is the next step.
- ◆ **Implementation**—Code is put into production.
- ◆ **Maintenance**—As errors are found or enhancements required, code is modified, tested, and placed into production.
- ◆ **Disposal**—At some point, legacy code is retired because the system is no longer needed or has been replaced by completely new systems.

Spiral Model

The spiral model starts in the middle with the conceptual model of what must be done and spirals outward through its phases, which repeat, at ever widening paths. One approach to the spiral model is

1. Develop a preliminary design.
2. Develop a prototype from the design.
3. Develop the next prototype.
4. Evaluate.
5. Define further requirements.
6. Plan and design another prototype.

7. Construct and test this prototype.
8. Repeat steps 3–7 until the customer is satisfied that the prototype meets the requirements.
9. Construct the system.
10. Thoroughly test the final system.

An additional model is the spiral model constructed like the waterfall model with the element of risk analysis added. This model is credited to Barry Boehm, chief engineer at TRW in 1988. In essence, four operations are repeated until the right design is created, which is then put into production. The four operations are

- ◆ **Planning/review**—Determine the objectives of the system to be developed.
- ◆ **Risk analysis, prototype**—First, identify all alternative solutions and perform a risk analysis. Resolve the risks and create the prototype.
- ◆ **Engineering**—Develop and verify the product requirements. Validate the design. Do a detail design and validate it. Code a test product.
- ◆ **Plan the next phase**—Review for customer satisfaction. Do requirements planning, development planning, and integration planning, and create a test plan.

Rapid Application Development

Rapid application development (RAD) recognizes that the result of software development is a product that meets economic, reliability, and speed-of-development goals. It seeks to develop a product that has 80% of what is desired but is produced in 20% of the time normally required to meet 100% of the goals. A common saying is that a RAD project has a strong chance of developing the product in the timeframe desired if the company is willing to sacrifice either economy or quality. And,

that it has a better chance of achieving its goal if the customer is willing to sacrifice both economy and quality.

Security Control Architecture

Security control architecture consists of the following:

- ◆ **Process isolation**—The capability to run different processes and separate them from one another.
- ◆ **Hardware segmentation**—The isolation of software processes and data via the segmentation of hardware.
- ◆ **Memory protection**—Virtual memory is divided into segments. Each process uses its own segment, and the system keeps its own internal processing separate from that of user mode processing (the running of applications).
- ◆ **Least privilege**—Processes have no more privileges than necessary to perform functions.
- ◆ **Separation of duties**—It is possible to assign privileges on the system so that related privileges are segregated—for example, backup and restore.
- ◆ **Layering**—A structured, hierarchical design of system function. Layers communicate through calls via defined interface.
- ◆ **Security kernel**—Hardware, firmware, and software that implement a reference-monitor concept.
- ◆ **Modes of operation**—Different system uses are separated into privileged and unprivileged.
- ◆ **Accountability**—With one user per account, you must be able to identify the individual's activity on a system.

Software Development Methodologies

Good software can be developed using many methodologies. Some methodologies can be performed only with certain programming languages. The following major development methodologies are in use today:

- ◆ **Structured programming**—Requires the programmer to be aware of the flow and control of the program. Structured programming is based on the principals of
 - Modularity
 - Top-down design
 - Limited control structures
 - Limited scope of variables
- ◆ **Object-oriented programming**—The emphasis is on describing the object and its data, methods, and interface.
- ◆ **Computer-aided software engineering (CASE)**—Computer applications that are designed to assist program development.
- ◆ **Prototyping**—A quick model of the program is made and viewed by users; then it's remodeled until it is approved. Then the working program is made.

Coding for Security

Ways to improve software include

- ◆ Eliminate buffer overflows.
- ◆ Prevent array indexing errors.
- ◆ Use access control.

DOMAIN 5, "CRYPTOGRAPHY"

Uses of Cryptography

Cryptography can be used for many purposes. The following are several primary uses:

- ◆ **Confidentiality**—Preventing, detecting, or deterring unauthorized access to information. Confidentiality of information can be obtained through both symmetric and asymmetric encryption.
- ◆ **Integrity**—Preventing, verifying, and detecting the alteration of data or information you have sent. Hash algorithms are typically used to provide for integrity of information.
- ◆ **Authentication**—Identifying an individual or verifying that the individual is part of a certain group. You typically can authenticate someone based on one of three attributes:
 - Something the person knows, such as a password
 - Something the person has, such as a token
 - Something the person is, or biometrics

Encryption is used by all three authentication methods.

Nonrepudiation is critical when it comes to digital signatures. It deals with proving in a court of law that someone was the originator. Nonrepudiation is a feature of asymmetric encryption that allows you to prove that someone actually sent a message. It is equivalent to an actual signature.

Cryptographic Methods and Algorithms

To understand cryptography better, you should make sure you know the basic definitions and the difference between symmetric and asymmetric algorithms, MACs, hash functions, and other cryptographic basics.

Definitions

Many cryptographic discussions assume knowledge of these basic definitions:

- ◆ **Plaintext**—A message in its original form. Remember that any type of message can be encrypted. So, even though the word has *text* in its name, plaintext is really a generic term and can refer to an executable, a zipped file, a word-processor document, a spreadsheet, or any type of information you would want to keep protected and secure. This is the data before anything has been done to it.
- ◆ **Ciphertext**—A message after it has been encrypted.
- ◆ **Encryption**—The process of taking a plaintext message and converting it to ciphertext.
- ◆ **Decryption**—The process of taking ciphertext and converting it back to a plaintext message. The key thing with encryption and decryption is this: If you take a plaintext message, convert it to ciphertext, and then decrypt it back to plaintext, the plaintext, decrypted message must match the original plaintext message that was inputted into the encryption algorithm.

Symmetric

Symmetric encryption is often called *single-key* or *secret-key encryption* because a single key is used for both encryption and decryption of the information.

This is one problem with symmetric-key encryption: The key must be sent over a secure channel. The other problem with symmetric key encryption is nonrepudiation. If we are both using the same key, how can one of us prove in a court of law that the other one sent the message? DES (data encryption standard) and triple DES are the most popular symmetric key encryption schemes used.

Asymmetric

Asymmetric encryption is often called *two-key encryption* or *public-key encryption*. It involves two keys: a public and a private key. The public key is given to anyone who wants it, and the private key is kept secret by the user. Anything that is encrypted with one key can only be decrypted with the other key.

If asymmetric encryption is so powerful, why do you need symmetric encryption? The reason is speed. RSA is the asymmetric algorithm of choice and is used in most implementations that utilize this type of encryption.

MACs

Message authentication codes (MACs) are used to ensure the message has not changed in transit and therefore protect it against integrity attacks.

Hash Function

A *hash function* is a one-way transformation that cannot be reversed.

Digital Signature

Digital signatures are used to ensure nonrepudiation.

Encryption Facts

The longer the key, the more possible potential values for the key, which means it will take longer to guess.

A rule of thumb is that the usefulness of the information should be less than the time it takes to brute-force the encryption.

A *one-time cipher* is often considered to be unbreakable encryption. That is not really a completely accurate statement. The reason people make this claim is that each time you encrypt a message, you use a new key. So, you would never ever use the same key twice.

PKI

Using asymmetric or symmetric encryption, you need to have keys in order to encrypt or decrypt the information. To communicate with a couple of people, managing keys yourself is easy, but what happens when you roll out encryption across a large enterprise? This central server is called a *public key infrastructure (PKI)* server.

Attacks Against Encryption

Encryption is not a foolproof answer—encryption algorithms can also be attacked. Attacks fall into two categories, general and specific.

General Attacks

Four general attacks can be performed against encrypted information:

- ◆ **Ciphertext only**—With a ciphertext-only attack, the only thing the cryptanalyst has is encrypted text.

- ◆ **Known plaintext**—Known-plaintext attacks imply that for a given message the cryptanalyst somehow was able to find the original plaintext message that was used to generate the ciphertext.
- ◆ **Chosen plaintext**—In some cases, access to the device that generates the encryption can be obtained without obtaining the key. In this case, you could feed in whatever plaintext you want and receive the corresponding ciphertext. This is one step easier than the known-plaintext attack.
- ◆ **Chosen ciphertext**—The last, general attack is a very sophisticated attack. In this attack, you can pick the ciphertext and the system will give you the corresponding plaintext.

Specific Attacks

Each general attack type is defined by whether ciphertext or plain text is available for use, but specific attacks can be defined by their methodologies:

- ◆ **Brute-force**—Because the goal is to find the key, you could try every possible combination.
- ◆ **Replay attack**—Involves taking encrypted information and playing it back at a later point in time.
- ◆ **Man-in-the-middle attack**—The attacker has inserted herself in the middle of the communication.
- ◆ **Meet-in-the-middle attack**—A potential vulnerability that exists with double DES, it is the reason double DES is not used.
- ◆ **Birthday attack**—A birthday attack against hash functions deals with trying to find two different messages that hash to the same value.

DOMAIN 6, “SECURITY ARCHITECTURE AND MODELS”

Examining the Differences Between Government and Industry Models

Historically, government computer security issues have centered on confidentiality—making sure unauthorized individuals cannot access information. On the public (or commercial) side, concerns have been of the correctness or integrity/consistency of data.

Security Models

Security models are attempts at organizing the management of security in an environment. Other models, discussed in other chapters, are examined here for comparison.

Clark-Wilson

The Clark-Wilson model emphasizes data integrity and does so for commercial activities. It uses software engineering concepts such as abstract data types, separation of privilege, allocation of least privilege, and nondiscretionary access control. Clark-Wilson has three integrity goals:

- ◆ Prevent unauthorized users from making modifications
- ◆ Prevent authorized users from making improper modifications
- ◆ Maintain internal and external consistency

Access Control Lists

In the Access Control List (ACL) model, objects (the resources) are assigned lists of approved subjects (users and groups). Each entry in the list consists of user identification of some form and the approved access level. Access levels are appropriate for the resource—hence for files, levels can be read, write, read/write, and so on, whereas for printers, levels can be manage or print. Subjects, the users and groups, are assigned some kind of identification.

Comparison of Common Security Models for Access Control

Table 3 compares the security models for access control.

TABLE 3
SECURITY MODELS FOR ACCESS CONTROL

<i>Name of Model</i>	<i>Government Model</i>	<i>Primary Directive</i>
Biba	Yes	Confidentiality
Bell-LaPadula	Yes	Confidentiality
Clark-Wilson	Yes	Integrity
Access Control Lists	No	Attempts at both confidentiality and integrity but limited to proper application

Security Architecture

A security architecture is the sum of the components used and the way they are put together to build security functionality into a computer operating system, device, or system.

Open System Versus Closed System

Table 4 compares open and closed systems.

TABLE 4
AN OPEN SYSTEM VERSUS A CLOSED SYSTEM

<i>System Item</i>	<i>Open</i>	<i>Closed</i>
User interface	Standard	Nonstandard
User access to system	Total	Limited to a single application or language

Security Principles

Some security principles to understand are

- ◆ **Trusted Computing Base (TCB)**—The sum of the security functions of the system.
- ◆ **Execution domain**—The OS system area is protected from tampering and accidental modification. Another layer, the user area, is set aside for application programs.
- ◆ **Layering**—Processes do not do everything. Processes are layered, with each layer having a specific job.
- ◆ **Abstraction**—Acceptable operations are characterized, not spelled out in detail.
- ◆ **Process isolation**—Many processes can be running without interfering with each other.
- ◆ **Least privilege**—A process has only the rights and access it needs to run; only processes that need complete privileges run in the kernel, and other processes call on these privileged processes only as needed.
- ◆ **Resource access control**—Access to resources is limited.

- ◆ **Security perimeter**—The boundary of the TCB. A security kernel and other security-realized functions operate within this perimeter. A security kernel is the implementation of the reference monitor concept.
- ◆ **Security policy enforcement**—The policy set for the system must be operational for the system to be operational. The security policy is always followed.
- ◆ **Domain separation**—The objects that a subject can access become its domain. The user doesn't need to access the security kernel, for example, so the domain of the TCB is separated from that of the user.

Security Modes

Security modes are indications of the currently operating function of a system. They are

- ◆ **Dedicated**—No restrictions. All users can access all data. All users have clearance for all data on the system and have signed nondisclosure agreements for all information stored and processed. The users have a valid need to know for all information.
- ◆ **System high**—All users have access approval and clearance for all information on the system. Users have clearance for all information. They have a need to know for some of the information and have signed nondisclosure agreements that require them not to share the information.
- ◆ **Compartmented**—Users have valid clearance for most restricted information processed on the system, formal access and nondisclosure for that information, and need to know for that information. Data is partitioned. Each area of data has different requirements for access. Users of the system must meet the requirement for the area they wish to access.

- ◆ **Multilevel secure (MLS)**—Users have different levels of clearance to different levels of information (think Bell-LaPadula). Some do not have valid personnel clearance for all information. All have valid need to know for that information to which they have access.
- ◆ **Controlled mode**—Multilevel access in which a more limited amount of trust is placed in the hardware/software base of the system. This results in more restrictions on classification levels and clearance levels.
- ◆ **Limited access mode**—Minimum user clearance is not cleared, and maximum data sensitivity is not classified by sensitivity.

Covert Channels

A *covert channel* allows an object with legitimate access to information to transfer the information in a manner that violates the system security policy. Two types of covert channels exist—covert storage channels and covert timing channels.

Information Security Standards

Standards for information security exist at national and international levels. The most commonly known and followed are as follows:

- ◆ Orange Book—Trusted Computer System Evaluation Criteria (TSEC), 1985
- ◆ UK Confidence Levels, 1989
- ◆ ITSEC (1991) Information Technology Security Evaluation Criteria (from the German and French Criteria, the Netherlands, and the United Kingdom)
- ◆ Canadian Criteria, 1993, Canadian Trusted Computer Product Evaluation Criteria (CTCPEC), a combination of ITSEC and TCSEC
- ◆ Federal Criteria, 1993 (draft Federal Criteria for Information Technology Security); later merged into Common Criteria

Orange Book

The certification emphasis of the Orange Book is confidentiality. The concept of a secure, or trusted, system is divided into a series of classifications that range from minimal protection to verified protection.

The Orange Book outlines the evaluation criteria and gives an objective measure for acquisition. It divides operating systems into four primary divisions around three different concepts. The concepts are

- ◆ Ability to separate users and data
- ◆ Granularity of access control
- ◆ Trust or overall assurance of the system

The primary divisions are

- ◆ **D**—Minimal protection
- ◆ **C**—Discretionary protection
- ◆ **B**—Mandatory protection
- ◆ **A**—Verified protection

Table 5 lists and describes the Orange book classifications.

600 FAST FACTS

TABLE 5
ORANGE BOOK CLASSIFICATION

<i>Class</i>	<i>Title</i>	<i>Description</i>
<i>D: Minimal protection: Have been evaluated but don't meet standards for other classes</i>		
<i>C: Discretionary protection: Need to know protection, accountability of subjects, accountability of actions, and audit</i>		
C1	Discretionary security protection	Separation of users and data; enforces access limitations; users use data at the same level of security
C2	Controlled access protection	More granular; user is more individually accountable; logical procedures, auditing, and resource isolation; security policy enforcement; accountability, assurance; controls who can log in; access to resources is based on wishes of users; log of user actions
<i>B: Mandatory protection: Integration of sensitivity labels, labels used to enforce mandatory access rules, specification of TCB, reference monitor concept implemented</i>		
B1	Labeled security protection	Accurate labeling of exported information
B2	Structured protection	Formal security model; discretionary and mandatory access control extended to all subjects and objects; covert channels are addressed; TCB has protection-critical and nonprotection-critical elements; trusted facility management (systems admins and operator functions and configuration management control); system is relatively resistant to penetration
B3	Security domains	Reference monitor must mediate all access of subjects by objects and is tamper-proof; unauthorized code is excluded; security policy enforcement; complexity minimized; security administrator supported; audit expanded; and system recovery required; system is highly resistant to tampering
<i>A: Verified Protection</i>		
A1	Verified design	Functionally equivalent to B3, but verification techniques are used against the formal security policy; can give high degree of assurance; TCB is correctly implemented

Information Technology Security Evaluation Criteria

This European standard was developed in 1991 by Germany, France, the Netherlands, and the United Kingdom. In 1998, Finland, France, Germany, Greece, Italy, the Netherlands, Norway, Spain, Sweden, Switzerland, and the United Kingdom agreed to recognize ITSEC certificates from qualifying certification bodies.

Differences Between the Orange Book and ITSEC

ITSEC incorporates many of the items first expressed in the Orange book; however, differences such as the following do exist:

- ◆ Unlike the Orange Book, which concentrates on confidentiality, ITSEC addresses the triple threat of loss of confidentiality, loss of integrity, and loss of availability.

- ◆ In the specifications, the Target of Evaluation (TOE) is the product or system to be evaluated. In ITSEC, the TOE's functionality (can it provide this security function) and assurance (how do you know it is providing this functionality) are evaluated separately.
- ◆ ITSEC does not require the security components of a system to be isolated into a TCB.
- ◆ ITSEC provides for the maintenance of TOE evaluation. Some systems can maintain certification after patches without formal reevaluation.

Table 6 lists and describes the ITSEC levels of evaluation.

TABLE 6
ITSEC LEVELS OF EVALUATION

<i>Level</i>	<i>Description</i>
EO	Inadequate
E1	Definition of security target and informal architecture design exists. User/Admin documentation on TOE security exists. TOE is uniquely identified, and documentation exists that includes delivery, configuration, startup, and operations. The evaluator tests the security functions. Secure distribution methods are utilized.
E2	Informal, detailed design and test documentation are produced. Separation of TOE into security enforcing and other components. Audit trail of startup and output is required. Assessment includes configuration control, developer's security, and penetration testing for errors.
E3	Source code or hardware drawings must accompany the product, and a correspondence between design and source code must be shown. Standard, recognized implementation languages are used. Retesting is required after correction for errors.
E4	Formal security model. Semiformal specification for security enforcing functions, architecture, and detailed design. Sufficient testing. TOE and tools are under configuration control. Changes are audited, and compiler options are documented. TOE retains security after a restart from failure.

<i>Level</i>	<i>Description</i>
E5	Relationships between security enforcing components are defined in architectural design. Integration processes and runtime libraries are provided. Configuration control is possible independently of the developer. Configured security enforcing or relevant items can be identified. There is support for variable relationships between them.
E6	Formal description of architecture and security enforcing functions with correspondence between formal specification through source code and tests. All TOE configurations are defined in terms of the architecture design, and all tools can be controlled.

Common Criteria

The "Arrangement on the Recognition of Common Criteria Certificates in the Field of IT Security" was signed as a mutual recognition arrangement in 1998 by government organizations from the United States, Canada, France, Germany, and the United Kingdom. This international standard, known as Common Criteria, has the following as its objectives:

- ◆ Ensure IT product evaluations are performed to high and consistent standards.
- ◆ Guarantee that evaluations contribute to the confidence in the security of the products.
- ◆ Increase the availability of evaluated, security-enhanced IT products.
- ◆ Eliminate duplicate evaluation.
- ◆ Continuously improve efficiency and cost-effectiveness of security evaluations and certification/validation process for IT products and protection profiles.

Common Criteria Evaluation Assurance Packages or Levels

EALs are combinations of assurance components. EAL7 is the highest with international recognition:

- ◆ **EAL1**—Functionally tested. Confidence in correct operation is required, but threats are not serious. Due care has been exercised with respect to protection.
- ◆ **EAL2**—Structurally tested. Delivery of design information and test results are consistent with good commercial practice. Low to moderate level of independently assured security. Many legacy systems can be evaluated at this level.
- ◆ **EAL3**—Methodically tested and checked. Security engineering at design states; requires minimal alteration of existing sound development practices to meet. (Grey box testing, search for obvious vulnerabilities.)
- ◆ **EAL4**—Methodically designed, tested, and reviewed. Use of positive security engineering and good commercial development practices is rigorous but does not require substantial specialist knowledge, skills, or testing. Independent search made for obvious vulnerabilities.

- ◆ **EAL5**—Semiformally designed and tested. Semiformally tested using rigorous commercial development practices and application of specialized security engineering techniques. High level of independently assured security in planned development; rigorous developmental approach.
- ◆ **EAL6**—Semiformally verified, designed, and tested. Specialized security engineering techniques in rigorous development environment. Protection of high-value assets against significant risks. Modular, layered approach to design; structured presentation of the implementation. Independent search for vulnerabilities ensures resistance to penetration, systematic search for covert channels, development environment, and configuration management controls.
- ◆ **EAL7**—Formally verified, designed, and tested. This is used for extremely high-risk situations or high-value of assists. White box testing is used.

A Comparison of the Orange Book, ITSEC, and Common Criteria

Table 7 compares the Orange book, ITSEC, and Common Criteria.

TABLE 7
A COMPARISON OF THE ORANGE BOOK, ITSEC, AND COMMON CRITERIA

<i>Orange Book</i>	<i>TCSEC</i>	<i>ITSEC</i>	<i>Common Criteria Evaluation Assurance Level</i>
D	Minimal protection	E0	EAL0 EAL1
C1	Discretionary security protection (discretionary access control, identification and authentication, system architecture, system integrity, security testing, documentation)	F1+E1	EAL2
C2	Controlled access protection (object reuse and audit)	F2+E2	EAL3
B1	Labeled security protection (labeling, label integrity, design verification)	F3+E3	EAL4
B2	Structured protection (covert channel, device labels, subject sensitivity)	F4+E4	EAL5

<i>Orange Book</i>	<i>TCSEC</i>	<i>ITSEC</i>	<i>Common Criteria Evaluation Assurance Level</i>
	labels, trusted path, trusted facility management, configuration management)		
B3	Security domains (intrusion detection, security administrator role definition)	F5+E5	EAL6
A1	Verified design (verified design, more documented version of B, trusted distribution)	F6+E6	EAL7

Uses for IPsec

IPsec is not just for encryption; its many uses include

- ◆ **Access control**—Access can be restricted by identifying the IP address of the computer(s).
- ◆ **Connectionless integrity**—A checksum is calculated, and a hash is computed across the payload and is also encrypted.
- ◆ **Mutual computer authentication**—Prior to data transmission, each computer must authenticate to the other.
- ◆ **Confidentiality**—The information is protected during transit. If the information is captured, it cannot be easily interpreted because it is encrypted.
- ◆ **Data-origin authentication**—Each packet can be attributed to the sending computer.
- ◆ **Protection against replay attacks**—Three items, securities parameter index (SPI), sequence number, and IP address, identify each packet. If this tuple matches that of a previously received packet, IPsec considers this an attack and drops the packet.

DOMAIN 7, “OPERATIONS SECURITY”

Roles of Operations Security

Operations security can be used to do the following:

- ◆ Identify resources to be protected.
- ◆ Identify privileges to be restricted.
- ◆ Identify available controls and their types.

Resources to Be Protected

Operations security should be designed to protect the following:

- ◆ Computers, including servers, desktops, and laptops
- ◆ Routers, switches, and other networking appliances
- ◆ Printers
- ◆ Databases, including the database management software and content
- ◆ Security software and appliances (firewalls, intrusion detection systems, biometric devices, public key infrastructure)

604 FAST FACTS

- ◆ Media such as tapes, CD-ROMs, and disks
- ◆ Personal digital appliances (PDAs), phones, and wireless devices
- ◆ Modems and other communications devices
- ◆ Software, including licensed commercial software and custom applications
- ◆ Source code
- ◆ Documentation

Types of Controls

To fulfill its objectives, operations security uses many types of controls, such as

- ◆ **Operational controls**—These are day-to-day procedures, mechanisms that include physical and environmental protection, privileged entry commands, change control management, hardware controls, and input and output controls.
- ◆ **Audit and variance detection controls**—These are audit logs that contain information on the exercise of privilege and records of system activity.

- ◆ **Application software maintenance controls**—These controls monitor installation and updates to applications, and they keep a record of changes.
- ◆ **Technical controls**—These controls audit and journal integrity validations, such as checksums, authentication, and file system permissions.
- ◆ **Administrative or management controls**—These control personnel screening, separation of duties, rotation of duties, and least privilege.
- ◆ **Deterrent controls**—These controls reduce the likelihood of attack.
- ◆ **Preventative controls**—These controls protect vulnerabilities, reduce the impact of attacks, or prevent an attack's success.
- ◆ **Detective controls**—These controls detect an attack and can activate corrective controls or preventative controls.
- ◆ **Corrective controls**—These controls reduce the impact of an attack.

Table 8 lists and matches controls to types.

TABLE 8
SAMPLE CONTROLS MAPPED TO TYPES

<i>PC Control</i>	<i>Control Types from Different Schemas</i>	
Require passwords for access, require biometrics for authentication	Technical	Preventative
Disk locks	Technical	Preventative
Acceptable use policies, requiring virus check of portable media	Operational	Preventative
Checking for compliance	Audit and variance detection	Corrective
Using antiviral software	Technical	Preventative
Requiring file encryption	Technical	Preventative
Training in controls	Management	Preventative
Requiring that help desk or IT staff configure PCs, not users	Management	Preventative
Software code audit looking for buffer overflows	Technical	Input, output
Loading a personal firewall/IDS system	Technical	Detective

Role of Auditing Monitoring

Auditing, whether with logs or special intrusion detection, devices can be used to

- ◆ Audit for compliance to security policy.
- ◆ Audit for evidence of intrusion, attack, or compromise.

Intrusion Detection

Intrusion detection is accomplished by extracting data and by the recognition of traffic and traffic patterns.

A network-based IDS analyzes all traffic on the network. A central management station usually manages the information gathered by the host and network IDSs.

A host-based IDS requires loading software on the host machine. The software listens to traffic coming and going to and from its host machine. It can also take advantage of information in the computer's logs and monitor the integrity of the file system for a broader picture of changes and attempted changes.

One of the tuning mechanisms is the capability to set the number of errors or instances of unusual activity that will cause an alarm. This is called *setting the clipping level*.

Penetration Testing Techniques

To do a penetration test, you should do the following:

- ◆ Determine the target.
- ◆ Footprint or profile.
- ◆ Enumerate the network.
- ◆ Scan and enumerate services on the network.
- ◆ Operating system enumeration.
- ◆ Attack against a particular machine.

Countermeasures to Threats

Table 9 gives examples of common threats.

TABLE 9
COMMON THREATS WITH EXAMPLES

<i>Threat</i>	<i>Notes</i>	<i>Example</i>
Errors	Incorrect configuration.	Default, well-known passwords are not changed.
Omission	Patches are not applied.	Patches for IIS were not applied, and many IIS servers were infected with Code Red.
Fraud	Company assets are obtained by misrepresentation or modification of information.	Paycheck amounts were increased by claiming overtime hours not worked, customer records were stolen, or software was taken by employees for home use.
Misuse of information	Sensitive, private information is used for personal gain.	Earnings knowledge used to buy or sell shares (insider trading).
Employee sabotage	Employee uses knowledge of company operations and systems to destroy or damage assets.	Time-bombed code is loaded on servers by an administrator and destroys data the day after the employee is fired.

continues

606 FAST FACTS

TABLE 9 *continued*
COMMON THREATS WITH EXAMPLES

<i>Threat</i>	<i>Notes</i>	<i>Example</i>
Ignoring policy	Employees know the rules but do not obey them.	Accidents are caused by not following safety rules. Accidental destruction of data backup caused by leaving tapes in the trunk of a parked car during a summer heat wave when policy states immediate transport in air-conditioned vehicle.
Physical accidents	These are as a result of physical circumstances as opposed to system malfunction or inadvertent misuse of the system.	Electric shock or moving parts of printers.
Software malfunction	Bugs or security vulnerabilities	Buffer overflow causes a reboot or leaves the system open to compromise.
Loss of resources	Destruction of data center in full or in part.	Fire, flood, storm, bomb, or explosion.
Loss of infrastructure	Malfunction of equipment.	A router or switch dies.
Hackers and crackers	Attack on systems.	Loss of data, loss of reputation, and destruction of systems.
Espionage	Spies from another company join yours or pay your employees to provide internal information.	Soft drink formula is stolen from database by employee and sold to competition.
Malicious code	Code is run on a system with undesirable results.	Code Red, Nimda, I Love You, and so forth.

Employee Job Duties and Risks to Systems

Table 10 examines job duties and the risks they can pose to information systems.

TABLE 10 **JOB DUTIES AND THEIR RISKS**

<i>Job</i>	<i>Description</i>	<i>Access Level</i>	<i>Risk</i>
Computer operator	Do backups; run jobs; mount tapes; load paper in printers; record and report problems; operate devices, software products, system performance metering, heat control, and humidity controls	Console, tape/disk drives, printers, operations documentation, problem/change management system	Gains access to production data files, production maintenance and job control, program documentation; turns off logging (can lose audit trail); potential loss of system records due to not enough room on media
Operations analyst	Analyzes computer memory and hardware requirements, estimates use of disk and tape performance, advises on operations documentation, establishes backup recovery procedures, monitors service-level agreements, installs new hardware and telecommunications, replaces obsolete items, and troubleshoots	Test files, operation documentation, system performance reports	Access to production data files and production application programs

<i>Job</i>	<i>Description</i>	<i>Access Level</i>	<i>Risk</i>
Job control analyst	Job control language, assists application programmers, reviews production problems using problem change management process, tests and implements new features, and assists in product troubleshooting	Test job control files, job scheduling files, operations documentation, problem/change management system	Access to production data files, application programs, and job control files
Production scheduler	Plans, creates, and coordinates computer processing schedules for production jobs and job streams; consults with end users and application programmers concerning production schedules; completes ad hoc jobs; reviews results in comparison to planned schedules; and updates and issues monthly billing schedules	Job scheduling files, operations documentation, problem/change management system	Access to production files, data files, production application programs, and job control files
Production control analyst	Prints, balances, and distributes reports and records; manages printer, burster, and decollator; balances required reports; assists production scheduler; and performs inventory counts and computer supplies	Computer equipment, supplies and reports, and problem/change management system	Delivers reports to wrong individuals, theft of supplies
Tape librarian	Collects input tapes; sends/receives tapes from offsite storage; maintains tapes and cartridges; ensures adequate supply, tape storage, and vault; ensures critical backup; pulls historical files and stores at local tape vault or ships to offsite location; maintains logs; and controls physical inventory tape library	Automated tape library, problem/change management system	Production data files, application programs, and job control files

Countermeasures to employee risks include

- ◆ Provide clear definition of authority.
- ◆ Structure along functional lines.
- ◆ Ensure that any type of fraudulent behavior requires the collaboration of two or more individuals.
- ◆ Separate job functions when combining them provides too much control.
- ◆ Rotate people within their own areas.
- ◆ Prevent family members from holding jobs in areas you would not combine into one person's responsibilities.
- ◆ Provide clean, accurate, detailed job descriptions.
- ◆ Include as part of every employee performance review, evaluation, and consideration for raise and promotion the employee's observance of security practices.
- ◆ Provide annual training for all employees.
- ◆ Encourage IT security to work with other security specialists, such as plant and physical security.
- ◆ Maintain a standards manual, and enforce the standards.
- ◆ Require vacations be taken, and require that they be taken contiguously.
- ◆ Require sophisticated access controls at the entrances to sensitive areas and systems.

608 FAST FACTS

Countermeasures to Internet threats include

- ◆ **Footprinting/enumerating the network**—Most information gained here is public knowledge. You can, however, obscure some information.
- ◆ **Scanning/enumerating services**—Block all unnecessary inbound and outbound ports.
- ◆ **OS enumeration**—Because many operating system identity hints or direct identification information are returned in banners (notices returned when inquiries are made), where possible change or eliminate the banner presented by services.
- ◆ **Penetration testing**—Become knowledgeable of the tools and tests hackers use. Develop or find tools that are countermeasures to these tools and methods.

Countermeasures to physical threats include

- ◆ Don't build near explosion hazards, and don't locate a data center near any explosives. In addition, diesel-powered generators should not be located near the data center.
- ◆ To avoid windstorm damage, don't have exterior windows and provide protection from possible falling trees or manmade structures such as towers.
- ◆ Don't place the data center on lower floors. Break-ins occur more often on lower floors.
- ◆ Do not externally label data center locations or advertise in it phone books, Web sites, and so forth.
- ◆ Avoid basement locations. Water damage can result from flooding. Use watertight seals and reroute pipes and conduits away from the data center if possible.
- ◆ Don't place media storage areas/vaults near flammable or explosive material or near compressors, water, and gas tanks.

- ◆ Subdivide rooms with firewalls or man traps, and keep fire doors closed.
- ◆ Use noncombustible building materials.
- ◆ Store paper media separately from equipment.

The Role of Administrative Management

Administrative management, the management of all things administrative, can serve a critical role in operations security. Managers must concern themselves with legal compliance, risk management, and fiduciary (monetary) responsibility. These are impacted by operations security. In addition, management plays a key role in promoting education on security, overseeing compliance, participating in policy-making and enforcement, ensuring cross-departmental involvement, and approving funding.

Principles of OPSEC

Least privilege, separation of duties, and change management can improve security and reduce the risk of fraud and accidental loss of data or data integrity. However, many other operations and best practices contribute to the stability and security of information. Some of them are discussed in other domains. Legal issues such as legal requirements; the standards of due care/due diligence; and record retention, privacy, and protection are discussed in the legal domain. Data backup is discussed in the Disaster Recovery and Business Continuity domain. Additional operations security concepts and best practices are

- ◆ Privileged operation functions
- ◆ Email security, including antivirus controls

- ◆ Protecting sensitive information and media
- ◆ Change management

Antiviral Controls

For antiviral controls to work, the following must be true:

- ◆ Antiviral products must be installed on servers and desktops.
- ◆ Automatic, regular updating of both engine and patterns is a must at the server and desktop levels.
- ◆ Server-side products should be configured to use additional features. Blocking of executable attachments to email is one example of a server-side feature.
- ◆ Attention should be paid to new viral/worm vectors. All infections will not come from email or desktop systems.

Management of Sensitive Data

Sensitive data must be managed in order to protect it. The following techniques will assist you in protecting data:

- ◆ **Creation**—All data, however it is obtained, should immediately be classified and labeled.
- ◆ **Handling**—All data within the data center must be properly handled to ensure viability and confidentiality. Protect media by keeping it in its original packaging and away from direct exposure to heat, sunlight, and electrical shock or damage from dropping.

- ◆ **Storage**—Provide environmental controls such as the ideal temperature, ideal humidity level, and freedom from dust and dirt.
- ◆ **Cleaning**—Wax and cleaning agents should not be used in the computer room or on storage area floors.
- ◆ **Destruction**—When it is no longer necessary to maintain data, the data should be destroyed. Common practices include clearing and purging.

Change Management Control

Computer operations should institute a change management control system for IT infrastructure. The first step in the process should be to develop detailed documentation on the following:

- ◆ Network configuration
- ◆ Computer configuration
- ◆ System parameters and settings
- ◆ Application configuration
- ◆ Device configuration
- ◆ Locations for all computers, devices, media storage, and other parts of the infrastructure
- ◆ Job titles and descriptions of duties
- ◆ Test environment specifications
- ◆ Disaster and continuity plans
- ◆ Other aspects of computer operations

DOMAIN 8, “BUSINESS CONTINUITY PLANNING AND DISASTER RECOVERY PLANNING”

Mandated Plans

Interagency Contingency Planning Regulation is a regulation that mandates that financial institutions in the U.S. will have a disaster recovery plan. It was developed by the Financial Institutions Examination Council.

Differences Between DRP and BCP

Disaster recovery is the process of bringing back into production a critical business process that has been crippled or destroyed by some catastrophic event. Disaster recovery planning is the process of developing a plan to do so. Business continuity planning seeks to minimize the impact of catastrophic events on critical business processes, get the processes up and operational should some event occur, and bring the company back to full recovery after the immediate crisis has passed.

Business Continuity Planning Process

The business continuity planning phases are

- ◆ Test plans.
- ◆ Implement plans.
- ◆ Maintain plans.

- ◆ Determine the scope of the plan.
- ◆ Perform business impact analysis.
- ◆ Develop operational plans for each business process.

Business Impact Assessment

A business impact assessment (BIA) is the process by which a business’s critical services are identified and a maximum tolerable downtime (MTD) for each is determined. The MTD, sometimes also known as the recovery time objective (RTO), is the timeframe within which the critical service must become operational to ensure the business will survive.

Operations Plan

An operations plan should include

- ◆ **Preventative measures**—Those operations that might prevent events, such as fire, or mitigate the effect of an event should it occur.
- ◆ **Emergency response**—Includes the actions taken immediately to avoid injury and loss of life.
- ◆ **Recovery**—The process of putting critical operations back into operation.
- ◆ **Return to normal operations**—Transitional activity that returns the business to normal operations.

Insurance

Some items that should be questioned when assessing insurance policies are as follows:

- ◆ The type of risk covered
- ◆ The type of property policy valuation
- ◆ The need for specific additional insurance

Two types of risk can be quantified in the policy. *Named perils* specifies that the cause of the loss must be enumerated. *All risks* specifies that all causes of loss that are not explicitly excluded in the policy are covered.

Testing the Plan

Several possible ways to test a plan are

- ◆ **Desk checking**—Reading through the plan and thinking how it would be used
- ◆ **Reviewing the plan for currency**—Examining the plan in light of new business processes, procedures, equipment, and interruption events
- ◆ **Performing full parallel system tests**—Testing backup equipment, software, data copies, and personnel at a hot site or alternative location
- ◆ **Running through scenarios and mock emergencies**—Having people respond by walking through their responsibilities as if it were a real emergency
- ◆ **Testing calls to contractors**—Finding out whether emergency personnel, facilities, and restoration specialists can be reached at any time of the day or night
- ◆ **Remote operations testing**—Moving employees to alternative sites and asking them to operate remotely
- ◆ **Switching to the mirror system or site**—Performing a fail-over to a data vault
- ◆ **Reviewing insurance**—Making sure coverage is up-to-date and team members are aware of the steps to follow to ensure the best result
- ◆ **Testing by departments or business process groups**

Maintenance

A full review of the plan requires that each business process be examined to see whether the plan adequately addresses the needs of the current systems, equipment, facilities, and people. Among the items to review are:

- ◆ Is the insurance plan up-to-date?
- ◆ Have new processes and equipment been added, and are they covered in the plan?
- ◆ Has team membership been adjusted to include or exclude changes in personnel?
- ◆ Is testing being done?
- ◆ Are there new types of events or changes in the likelihood of them occurring?
- ◆ Have mergers, acquisitions, or divestitures occurred, and has the plan been adjusted?

Disaster Recovery Planning

Disaster recovery planning consists of steps to preserve and recover data processing, prevent disasters from happening, backup considerations, and the determination of alternative sites.

Recovering Data Processing

The planning process for disaster recovery should include seven things:

- ◆ **The scope of the plan**—Including what is to be recovered and whether it's servers, data, or facilities.
- ◆ **Procedures that help to prevent disasters.**
- ◆ **A list of resources that need to be available**—Including an alternative site, equipment, data backups, personnel, and so on.

612 FAST FACTS

- ◆ **The backup strategy**—This ensures current data is available for restoration.
- ◆ **A to-do list for the emergency response process.**
- ◆ **Step-by-step instructions for implementing the plan**—This includes getting processes into operation.
- ◆ **Phone numbers of restoration and alternative sites**—Including business, home, off-hour numbers, cell, and other alternative numbers for locating your contacts at these companies.

Antidisaster Procedures

It's especially important that disaster recovery planning pay attention to techniques for preventing disasters. The following items should be considered:

- ◆ Locking hubs, routers, and switches in their own wiring closets instead of leaving them exposed in public areas or housed with public utility access points
- ◆ Limiting access to data centers, server rooms, and equipment closets
- ◆ Using approved fire-retardant materials in the construction of data centers
- ◆ Providing fire-extinguishing equipment and sprinkler systems where appropriate
- ◆ Performing background screening of employees
- ◆ Using antivirus products on gateways, servers, and desktops
- ◆ Using screening firewalls, routers, and so on at both egress and ingress points into networks

Backup Issues to Consider

When planning backup, consider the following:

- ◆ **Data backup**—Traditional copy to tape or other media.
- ◆ **Alternative sites**—Moving operations to other locations.
- ◆ **Data vaulting**—Either the transaction or the data file is transmitted to an alternative location in real-time. This can include the capability for a hot backup to immediately take over processing.
- ◆ **Co-location**—An exact copy, say of a Web or e-commerce site, is located at an alternative site or ISP. The co-located site is immediately ready to take over serving pages, accepting orders, and so on if a problem occurs at the main location.
- ◆ **Hardware backup**—Duplicate hardware is available either at the main site or alternative location, or both. It can immediately be put into service and the latest backup restored.
- ◆ **Hardware- or software-based redundant array of inexpensive disks**—Fault-tolerant disk systems provide duplication of data or the capability to recover data in the face of drive failure. Several techniques are used.
- ◆ **Fail-over clustering**—Multiple processors operate in a cluster and provide the capability to automatically switch from malfunctioning units to functioning units.

Alternative Sites

Different types of alternative sites can be selected. They include

- ◆ **Hot**—Completely configured with equipment, systems software, and appropriate environment.

- ◆ **Warm**—Partially configured with the possibility of having peripheral equipment such as printers.
- ◆ **Cold**—Only the basic environment (wiring, power, air conditioning, and so on) is available.
- ◆ **Redundant**—It's set up exactly like the primary site.
- ◆ **Mobile**—A site configured in a trailer or van; it can be operational anywhere. It's often brought to the company to be used while the primary site is being repaired.
- ◆ **Hybrid**—It's some combination of these types of sites.

DOMAIN 9, “LAW, INVESTIGATION, AND ETHICS”

Criminal, Civil, and Administrative Law

Criminal laws authorize the government to punish wrongdoers with financial penalties and incarceration. To convict a suspect under criminal law, the government must meet a high standard of proof—*proof beyond a reasonable doubt*—that the suspect intentionally did something wrong.

Civil laws, on the other hand, enable private parties to enforce their rights—such as contract, tort, and property rights—through court orders and monetary awards for damages.

Administrative law allows government agencies to interpret the laws they administer through official statements or regulations and to enforce those laws through investigations, fines, and other sanctions.

Intellectual Property Law

The major categories of intellectual property law available are

- ◆ **Patents**—A patent grants to its owner the exclusive right to make, use, or sell an invention covered by the patent. A patent can cover a physical invention or a business process, such as a unique process executed by software. To obtain a patent, an inventor must apply to the U.S. Patent and Trademark Office (USPTO). Often, the inventor must wait two or three years before the USPTO decides whether to grant the patent.
- ◆ **Copyrights**—Copyright law grants to the owner of a copyright the exclusive right to copy and make derivative works from the copyrighted material. Copyright covers expressions of ideas, such as written words, pictures, sounds, software code, and even live performances. But copyright covers only the expressions of the ideas, not the ideas themselves.
- ◆ **Trade secrets**—Trade secret law allows the owner of a trade secret to prevent others from using or exploiting the secret. A trade secret might be something like a customer list or an algorithm for searching through data on a network. Trade secret law applies automatically to information a company treats as a trade secret.

Sales and Licensing

When a programmer or contractor is hired to write software, the employer typically obtains an agreement that all the programmer's or contractor's work product (inventions, copyrights, and trade secrets) are sold and assigned to the employer. This arrangement is known as *work for hire*.

614 FAST FACTS

A *license* is typically a contract that allows each customer to use the software (and the patents, copyrights, and trade secrets therein) under restricted terms but does not allow the customer to remarket the software as his own. A license typically means a right to use but not to own.

Privacy

The United States has no comprehensive national law on privacy. U.S. privacy laws tend to apply on a sector-by-sector basis. Several laws that affect the use and protection of information systems and the data they manage are

- ◆ State laws and the federal Healthcare Insurance Portability and Accountability Act (HIPAA) generally require healthcare providers to maintain the confidentiality of patient information.
- ◆ The federal Gramm-Leach-Bliley Financial Modernization Act requires financial institutions to give customers notice about how their private information will be protected or shared with third parties.
- ◆ The Privacy Act limits the ability of federal government agencies to disclose to the public or other agencies information they have about individual citizens.
- ◆ Generally, no American law requires that companies post privacy policies with respect to people who visit their Web sites. However, many companies do elect to post privacy policies to make visitors feel more comfortable.

Generally speaking, employees have no right to privacy when communicating through corporate information resources if the employees are informed in advance that they have no privacy. Therefore, many corporations publish notices to employees to the effect that management might monitor their email or other electronic communications.

In contrast to the U.S., the European Union (EU) has more comprehensive rules on individual privacy. Traditionally, these rules have included restrictions on “transborder data flows” that would allow private data to flow to countries whose laws would not protect that data. The European Union’s Directive on Data Protection forbids the transfer of individually identifiable information to a country outside the EU unless the receiving country grants individuals adequate privacy protection.

To establish that data sent to the U.S. is granted adequate privacy protection, the EU and the U.S. government have negotiated a *safe harbor*. Under the safe harbor, participating U.S. companies voluntarily agree to protect personally identifiable information from the EU.

Federal Laws

Federal laws that impact information processing are

- ◆ The federal Foreign Corrupt Practices Act (FCPA) requires publicly owned companies to maintain adequate books and records and an adequate system of internal controls. Normally, the FCPA is enforced as administrative law by the U.S. Securities and Exchange Commission.
- ◆ The federal Gramm-Leach-Bliley Financial Modernization Act, and official guidelines published under the act, require financial institutions to implement a security program to safeguard private customer information in their possession.
- ◆ The U.S. Export Administration Regulations require that exporters obtain licenses before they export certain high-performance computers and microprocessors, as well as strong encryption. The U.S. Commerce Department’s Bureau of Export Administration (BXA) administers and enforces these export controls. Noncompliance can lead to administrative sanctions and criminal penalties.

Criminal Law

Criminal laws punish serious offenses against society. All criminal convictions, whether computer-related or otherwise, must rest on a particular preexisting law making the person's actions a crime.

The federal Computer Fraud and Abuse Act is a criminal law that punishes people who intentionally cause harm by accessing computers without authority.

The act generally forbids people from knowingly gaining unauthorized access to a computer of the U.S. government or a financial institution or a computer that is used for interstate or foreign commerce (which embraces many computers on the Internet), if that access leads to

- ◆ Classified or national security-related information
- ◆ Records of a financial institution
- ◆ Government records
- ◆ Information on a computer involved in interstate commerce
- ◆ An effect on the government's use of the computer
- ◆ Fraud
- ◆ Damage
- ◆ Trafficking in passwords
- ◆ Extortion

The federal Wiretap Act, 18 United States Code Section 2511, is a criminal law that punishes unauthorized interception of electronic communications in transit.

The key to an action being punishable as criminal is that the suspect *intentionally* do something wrong. Without intent to do something wrong, there can be no crime.

A banner warning that unauthorized access to a network is forbidden can help provide proof that a hacker intentionally committed a crime.

Computer Crime Investigation

The steps involved in the investigation of a crime are

1. Detect the intrusion.
2. Do whatever is necessary to avoid any additional damage and cut off the potential for liability, such as liability to trading partners who stand to be damaged by the incident.
3. Report the incident to management.
4. Conduct a preliminary investigation that includes assessing damage, witnesses, and whether a crime has occurred and determining what the investigation will need going forward.
5. Decide whether disclosure of the incident to government or the media is desired or required. It might be mandatory, for example, to disclose bank fraud to banking regulators.
6. Decide on a course of action, such as tightening of security, maintaining surveillance, or seeking prosecution.
7. Assign responsibility for conduct of the investigation, whether it is to internal staff, external consultants, or law enforcement. If a search warrant is required, law enforcement must show a court that probable cause exists to believe that a crime has been committed and a search/seizure is needed to investigate.
8. Pinpoint potential suspects (insiders, outsiders, or a conspiracy of both) and potential witnesses, and designate who should interview witnesses.

616 FAST FACTS

9. Plan and prepare for the seizure of target systems, including the possible need for special experts and a search warrant.
10. Designate a search and seizure team, including a lead investigator, an IT security specialist, a legal advisor, and technical staff.
11. Evaluate the risk to the target system before seizing it, including an anticipated reaction of the suspect and the risk that evidence will be destroyed.
12. Execute the seizure plan. Secure and search the location, preserve evidence, record each action (such as in a notebook), videotape the process, photograph the system configuration and monitor display, and move the system to a secure location.
13. Prepare a detailed report documenting facts and conclusions.

Evidence

Some evidence is stronger or more credible than other evidence. The credibility of evidence is usually determined by the trier of fact—in other words, the judge or jury in the court—based on the following:

- ◆ Strong evidence of a fact is called *direct* evidence; weaker evidence is called *circumstantial* evidence.
- ◆ To be *authentic*, evidence must be supported by something showing that the evidence is what it purports to be.
- ◆ The “hearsay rule” excludes from court a statement made outside the court that is repeated for the purpose of showing the statement is true.

- ◆ The “best evidence rule” says that to prove the terms of a “writing,” the original writing must be produced in court—not a copy—because the original is more reliable. When an electronic writing is at issue, you can most easily satisfy the best evidence rule with respect to that writing by persuading the court that the evidence being offered is an accurate representation of the writing.
- ◆ The *chain of evidence* is a series of records showing where evidence came from, who was responsible for it, what happened to it, how it was protected, whether it was changed, and so on.

Fourth Amendment to the U.S. Constitution

The Fourth Amendment to the U.S. Constitution protects citizens from unreasonable searches and seizures by government. Therefore, law enforcement normally needs a court-issued warrant before searching or seizing evidence, although there are exceptions, such as when evidence is in plain view.

Forensics

The techniques for seizing and preserving electronic evidence so as not to alter or destroy it are as follows:

- ◆ Restrict physical and remote access to the computer.
- ◆ If computer is off, do not turn it on.
- ◆ If computer is on, photograph the image showing on the screen and then unplug the computer.
- ◆ Do not touch the keyboard.
- ◆ Do all forensic analysis of the electronic evidence from a mirror copy of the disk on which the evidence is originally stored.

- ◆ Don't trust the subject computer's operating system; conduct analysis on a copy using the operating system of a trusted computer.

PC Examination Checklist

The steps in a computer forensics examination are as follows:

1. Before starting a computer forensics examination, get appropriate authority from corporate management. If the investigator is in law enforcement, a court-issued search warrant might be necessary.
2. If the machine is on, turn it off by pulling the plug. To record the state of the computer before it was unplugged, photograph the image displayed on the monitor.
3. Before moving the computer, document the hardware configuration with photographs and tags on cables. Collect, package, and label removable media such as floppy disks, tapes, and CDs present in the premises of the PC.
4. Transport the computer to a secure location.
5. Boot the computer without booting from the suspect hard drive itself. Boot from a floppy, or remove the hard drive and examine it using a separate computer dedicated to forensic examination.
6. Using forensic software, make a bit-stream image of the suspect drive; then run a hash of the suspect hard drive and the image to confirm the data in the two are the same. Next, document the system date and time. Forensics software can then be used on the image copy to run keyword searches through files, free space, and slack space.

Ethics

RFC 1087 declares unethical and unacceptable any activity which purposely

- ◆ Seeks to gain unauthorized access to the resources of the Internet
- ◆ Disrupts the intended use of the Internet
- ◆ Wastes resources (people, capacity, computers) through such actions
- ◆ Destroys the integrity of computer-based information
- ◆ Compromises the privacy of users

DOMAIN 10, "PHYSICAL SECURITY"

Elements of physical security are

- ◆ **Facility requirements**—Such as site selection and construction and perimeter control
- ◆ **Technical controls**—Such as card or token systems
- ◆ **Environmental/life and safety**—Such as power and fire issues
- ◆ **Physical security threats**—Such as weather and other natural events and intentional attacks
- ◆ **Elements of physical security**—Such as sensors and surveillance

Classification of Assets

Four physical asset classes are

- ◆ **Facility**—Building, rooms, workspace, backup storage area, and so on

618 FAST FACTS

- ◆ **Support**—Air conditioning, fire systems, electricity, communications, water, fuel supplies, and so on
- ◆ **Physical and components**—Hardware, including servers, printers, storage units, laptops, and workstations; desks; chairs; containers; and similar objects
- ◆ **Supplies and materials**—Disks and other removable media, paper supplies, waste material, and so on
- ◆ **Emergency support response**—The nearness of fire stations affects how great your fire risk is, for example.
- ◆ **Vulnerability to natural disasters**—Is the proposed location susceptible to earthquake, tornadoes, or hurricanes? Is it located below a dam? Is it in an approach path to an airport?
- ◆ **General building construction**—Building construction is a major topic in itself.
- ◆ **Computer room considerations**—The computer center should be a protected (point security) area within the building.

Countermeasure to Theft

Theft is controlled by the following:

- ◆ Authorizing (or hiring) trustworthy people
- ◆ Maintaining a corporate culture in which honesty is expected and normal
- ◆ Motivating people by good work environments and competitive remuneration
- ◆ Minimizing opportunities that would allow the easy theft of assets

Site Location and Construction

Site location and the construction of a building and the data center have an impact on the risks to systems. The following are some things to consider:

- ◆ **Vulnerability to crime, riots, and demonstrations**—Consider whether the location will make you vulnerable to such problems.
- ◆ **Adjacent buildings and businesses**—Do nearby business attract types of attention you don't want directed toward your information systems facility? If there is an adjacent building, can someone get from it into yours and, if so, is its security as strong as your own?

Physical Access Controls

Physical access control is essentially a perimeter control. You need to understand the following issues related to physical access controls:

- ◆ Perimeter control
- ◆ Access versus security tradeoff
- ◆ Response
- ◆ Doors
- ◆ Keys, including card systems and other tokens, and window construction

Doors and key are passive controls. More active measures require people or, in some cases, expensive automated measures such as a computer-controlled card-access system. The people could be guards or receptionists.

Power

Power issues and countermeasures are

- ◆ **Surges, spikes, and brownouts**—Use a UPS system, which provides power management and therefore provides an even source of supply to computer systems regardless of spikes and brownouts.

- ◆ **Outages**—Always prepare for outages by providing redundancy for all systems. This includes software, hardware, and processing.
- ◆ **Static**—In addition to temperature control, the control humidity is important because it can reduce or eliminate harmful static.

Environmental Controls

Demand adequate air conditioning and heating for computer systems.

Water Exposure Problems

Water exposure problems can be caused by things ranging from something as simple as a window open during a rainstorm to something as wide-ranging (and outside an individual organization's control) as a collapsed tunnel letting a river into most of downtown Chicago's sub-basement system. A short list of common problems includes

- ◆ **Flood**—Whether from weather or municipal facility problems.
- ◆ **Basements**—Water from an upper floor problem tends to result in flooded basements.
- ◆ **Roofs**—Leakage, burst drainpipes during heavy storms, and so on.
- ◆ **Snow load problems.**
- ◆ **Hurricane and other weather phenomena.**
- ◆ **Sprinklers.**
- ◆ **Air conditioning**—Often uses water as a coolant or heat transfer fluid.

Fire Prevention

Measures that can prevent fire or mitigate the damage it can cause are

- ◆ The materials used in a computer room should be as fireproof as practical.
- ◆ A fireproof media vault should be provided.
- ◆ Fire regulations should be known and observed by all employees.
- ◆ Fire drills.
- ◆ A no smoking policy.

Fire Extinguishers

Common fire extinguishers are Halon gas and carbon dioxide.

With the signing of the Montreal protocol in 1987, Canada, the United States, the European Community, and 23 other nations agreed to control the production and consumption of certain chlorofluorocarbon compounds (CFCs), including the Halon group. These ozone-depleting substances include some refrigerants and, relevant to this discussion, Halon 1211, Halon 1301, and Halon 2402. These Halons are used primarily in fire-extinguishing applications. The CFC compounds are implicated in the depletion of the ozone layer. The timetable for implementation of the Montreal protocols was advanced in 1992, and chlorofluorocarbon fire systems might not be a viable alternative for new, or even existing, installations. Halon systems are still used in special circumstances, but under severe regulation.

Regulations regarding the use of Halon vary, but typically include these recommendations:

- ◆ When planning fire protection for new installations, all alternative options (carbon dioxide, water, and so on) should be fully explored before deciding to use Halon.

620 FAST FACTS

- ◆ When Halon is used, full-discharge testing should be avoided in favor of alternative test procedures.

Tape and Media Retention Policy

Tape and other media should be stored in a protective environment, labeled, and their retention determined before they are stored. Other ways of caring for tape and media include

- ◆ **Restricted**—Storage areas need to be at least as carefully control as the area in which the data is used. All the access controls recommended for other restricted areas also are necessary in the media storage area.
- ◆ **Controlled**—Someone should have specific responsibility for keeping records of media entering the library and leaving it and for conducting frequent inventory of the contents. Any discrepancies should be followed up immediately.
- ◆ **Locked**—This is an elementary issue, but it is frequently ignored. Some form of an automatic locking mechanism is preferable so that carelessness cannot lead to a large exposure.
- ◆ **Protected from fire**—Media contain, as an acquired value, information that might be expensive or impossible to replace and that might be valuable to others. The storage area should be separated from the rest of the computer resource and should have its own independent fire protection. This could be elaborate in a large installation or fairly simple in a small shop.

Waste Disposal

Classified wastes should be handled as follows:

- ◆ Stored in separate containers.
- ◆ Collected frequently by security-cleared personnel.
- ◆ Retained in a secure area.
- ◆ Destroyed by cleared personnel, using an approved and effective method (shredding, incineration, and so on).
- ◆ Most personal computer operating systems do not actually *erase* data files when the operator says “erase” or “delete”; they set a flag indicating the file is “deleted.” The flag can be reset, and fragments of data might still exist. Programs exist specifically for the purpose of recovering deleted files. *Degaussing* is needed to ensure the erasure of data (a degausser generates a strong, varying magnetic field that randomizes the magnetic bits used to store data, thus data cannot be recovered).
- ◆ Data stored on most commonly available optical media (such as CD-ROM and DVD) cannot be erased; the medium must be destroyed thoroughly.
- ◆ Core dumps generated during program development (or sometimes when a program fails during operation) are sensitive waste. They contain a great deal of information that can possibly be accessed, and therefore should be destroyed, not just thrown out.
- ◆ Some kinds of computer memory stay “live” for a long time (up to years) even with the power turned off. An unauthorized user turning on the machine might get access to sensitive data.

These study and exam prep tips provide some general guidelines to help prepare for the CISSP exam. The information here is organized into three sections. The first section addresses pre-exam preparation activities and covers general study tips. Following this are some tips and hints for the actual test-taking situation. Before tackling those areas, however, you should think a little bit about how you learn.

LEARNING AS A PROCESS

To best understand the nature of preparation for the exams, it is important to understand learning as a process. You are probably aware of how you best learn new material. You might find that outlining works best for you, or you might be a visual learner who needs to “see” things. Whatever your learning style, test preparation takes place over time. Obviously, you can’t start studying for the CISSP exam the night before you take it; it is very important to understand that learning is a developmental process. And as part of that process, you need to focus on what you know and what you have yet to learn.

Learning takes place when you match new information to old. You have extensive experience in one or more domains of the CBK, and now you are preparing for the CISSP exam, which covers all 10 of them. Using this book, and supplementary materials, will not just add incrementally to what you know; as you study, you will actually change the organization of your knowledge as you integrate this new information into your existing knowledge base. This will lead you to a more comprehensive understanding of the domains and information security in general. Again, this happens as a repetitive process rather than as a singular event. If you keep this model of learning in mind as you prepare for the exam, you will make the best decisions concerning what to study and how much more studying you need to do.



Study and Exam Prep Tips

622 STUDY AND EXAM PREP TIPS

STUDY TIPS

There are many ways to approach studying, just as there are many types of material to study. The following tips, however, should work well for the type of material covered on the CISSP exam.

Study Strategies

Although individuals vary in the ways they learn, some basic principles apply to everyone. You should adopt some study strategies that take advantage of these principles. One of these principles is that learning can be broken into various depths. Recognition (of terms, for example) exemplifies a surface level of learning in which you rely on a prompt of some sort to elicit recall. Comprehension or understanding (of the concepts behind the terms, for example) represents a deeper level of learning. The ability to analyze a concept and apply your understanding of it in a new way represents an even deeper level of learning.

Your learning strategy should enable you to know the material at a level or two deeper than mere recognition. This will help you do well on the exam. You will know the material so thoroughly that you can easily handle the recognition-level types of questions used in multiple-choice testing. You will also be able to apply your knowledge to solve new problems.

Macro and Micro Study Strategies

One strategy that can lead to this deeper learning includes preparing an outline that covers all the objectives and subobjectives for the exam. You should delve a bit further into the material and include a level or two of detail beyond the stated objectives and subobjectives for the exam. Then, you should expand the outline by coming up with a statement of definition or a summary for each point in the outline.

An outline provides two approaches to studying. First, you can study the outline by focusing on the organization of the material. You can work your way through the points and subpoints of your outline, with the goal of learning how they relate to one another. You should be certain, for example, that you understand how each of the main objective areas is similar to and different from the others. Then, you should do the same thing with the subobjectives; be sure you know which subobjectives pertain to each objective area and how they relate to one another.

Next, you can work through the outline, focusing on learning the details. You should memorize and understand terms and their definitions, facts, rules and strategies, advantages and disadvantages, and so on. In this pass through the outline, you should attempt to learn detail rather than the big picture (the organizational information you worked on in the first pass through the outline).

Research has shown that attempting to assimilate both overall and detail types of information at the same time can interfere with the overall learning process. For the best exam performance, you should separate your studying into these two approaches.

Active Study Strategies

You should develop and actually exercise an active study strategy. You should write down and define objectives, terms, facts, and definitions. In human information-processing terms, writing forces you to engage in more active encoding of the information. Just reading over the information exemplifies more passive processing.

Next, you should determine whether you can apply the information you have learned by attempting to create examples and scenarios on your own: Think about how or where you could apply the concepts you are learning. Again, you should write down this information to process the facts and concepts in a more active fashion.

Common-Sense Strategies

Finally, you should follow common-sense practices when studying. You should study when you are alert, reduce or eliminate distractions, take breaks when you become fatigued, and so on.

Pretesting Yourself

Pretesting enables you to assess how well you are learning. One of the most important aspects of learning is meta-learning. *Meta-learning* has to do with realizing when you know something well or when you need to study some more. In other words, meta-learning is the ability to recognize how well or how poorly you have learned the material you are studying.

For most people, meta-learning can be difficult to assess objectively. Practice tests are useful in that they objectively reveal what you have learned and what you have not learned. You should practice test information to guide review and further study. Developmental learning takes place as you cycle through studying, assessing how well you have learned, reviewing, and assessing again until you think you are ready to take the exam.

You might have noticed the practice exam included in this book. You can use it as part of the learning process. The PrepLogic software on the CD-ROM also provides a variety of ways to test yourself before you take the actual exam. By using the practice exam, you can take an entire timed, practice test quite similar in nature to the actual CISSP exam. Although the CISSP exam is not electronic, the questions on the PrepLogic software are intended to simulate the type of questions you would find on the exam.

You should set a goal for your pretesting. A reasonable goal would be to score consistently in the 90% range.

See Appendix D, “Using the *PrepLogic Practice Tests, Preview Edition*,” for a more detailed explanation of the test engine.

EXAM PREP TIPS

The CISSP certification exam is a standardized, pencil-and-paper, proctored, multiple-choice, six-hour exam that reflects the 10 knowledge domains established by (ISC)².

The exam consists of 250 multiple-choice questions and a smaller number of “experimental” or “test” questions. The test questions are proposed new questions, and you are not penalized for answering them incorrectly, nor given extra points if you answer them correctly. This is the way new questions are tested for inclusion as part of the exam at a later date. The exam questions are not identified as being “experimental.”

The individual question booklets are printed in different orders to ensure that no two people sitting next to each other in the exam room have an exam created in the same order. If you take the exam more than once, you will see the same number of questions, but you won't see the exact same questions. This is because exam questions are periodically refreshed and the exam is given only at selected locations throughout the year.

Putting It All Together

Given all these pieces of information, the task now is to assemble a set of tips that will help you successfully tackle the CISSP certification exam.

624 STUDY AND EXAM PREP TIPS

More Exam Prep Tips

Generic exam preparation advice is always useful. Tips include the following:

- ◆ Pay particular attention to definitions.
- ◆ Review the current exam study guide and the “Process for Becoming a CISSP” guide on the (ISC)² Web site.
- ◆ Take any of the available practice tests. We recommend the ones included in this book and the ones you can create by using the PrepLogic software on the CD-ROM.
- ◆ Because there is a large amount of information to learn, it is tempting to spend time memorizing definitions. Remember that you need to be able to think your way through questions as well.

Tips for the Exam Session

The following generic exam taking advice you have heard for years applies when taking the CISSP exam:

- ◆ Take a deep breath and try to relax when you first sit down for the exam session. It is very important to control the stress you might (naturally) feel when taking exams.
- ◆ Carefully read all the information in the questions.
- ◆ Tackle the questions in the order in which they are presented. Skipping around will not build your confidence; the clock is always counting down.
- ◆ Do not rush, but also do not linger on difficult questions. The questions vary in degree of difficulty. Don't let yourself be flustered by a particularly difficult or verbose question.

- ◆ The exam is long. It can be helpful to make a rough calculation of how many minutes you can spend on each question and use this to pace yourself through the exam.
- ◆ Take advantage of the fact that you can return to and review skipped or previously answered questions. Record the questions you can't answer confidently, noting the relative difficulty of each question, on the scratch paper provided. After you have made it to the end of the exam, return to the troublesome questions.
- ◆ If session time remains after you have completed all questions (and if you aren't too fatigued!), review your answers. Pay particular attention to questions that seem to have a lot of detail.
- ◆ As for changing your answers, the general rule of thumb is don't! If you read a question carefully and completely and thought you knew the right answer, you probably did. Do not second-guess yourself. If, as you check your answers, one clearly stands out as being incorrectly marked, of course you should change it. If you are at all unsure, however, go with your first impression.

If you have done your studying and follow the preceding suggestions, you should do well. Good luck!

This exam consists of 250 questions reflecting the material you have covered in the chapters. These questions are representative of the types that you should expect to see on the actual exam.

The answers to all questions appear in their own section following the exam. We strongly suggest that when you take this exam, you treat it just as you would the actual exam at the test center. Time yourself, read the questions carefully, and answer all the questions to the best of your ability.

Most of the questions do not simply require you to recall facts but require deduction on your part to come up with the best answer. Some questions require you to identify the best course of action to take in a given situation. Run through the exam, and for questions you miss, review any material associated with them.



Practice Exam

626 PRACTICE EXAM

EXAM QUESTIONS

1. What does granting users access to objects under the principle of least privilege imply?
 - A. Full control
 - B. Minimal necessary access
 - C. No access
 - D. Role-based access
2. Which access control technique uses subject classification to determine access?
 - A. Discretionary access control
 - B. Access control lists
 - C. Mandatory access control
 - D. Rule-based access control
3. Which type of attack is a spoofing attack?
 - A. Monitoring attack
 - B. Spamming attack
 - C. Active attack
 - D. Passive attack
4. Which of the following are mechanisms of access control?
 - A. Physical, material, and discretionary controls
 - B. Administrative, logical, and physical controls
 - C. Administrative, supportive, and authentication controls
 - D. Confidentiality, integrity, and availability controls
5. A padded cell includes all but which of the following?
 - A. A simulated environment
 - B. Confidential data
 - C. Logging capabilities
 - D. Malicious action restrictions
6. An intrusion detection system (IDS) is which type of security measure?
 - A. Preventative
 - B. Reactive
 - C. Detective
 - D. Corrective
7. Intrusion detection systems are the weakest at identifying which of the following types of attacks?
 - A. Attempted unauthorized access to a secured object
 - B. Spoofing attacks
 - C. Denial-of-service attacks
 - D. Brute force attacks
8. What is the performance rating for biometric devices that is used to judge the relative effectiveness between similar devices from different vendors?
 - A. False rejection rate
 - B. False acceptance rate
 - C. Crossover error rate
 - D. Enrollment time
9. What is the most important aspect to consider when deploying a honeypot?
 - A. Logging
 - B. Legal ramifications
 - C. Protection of confidential data
 - D. Cross-platform support

10. The Bell-LaPadula security model was designed to address which of the following?
- A. Confidentiality
 - B. Integrity
 - C. Interoperability
 - D. Availability
11. The Biba security model was designed to address which of the following?
- A. Confidentiality
 - B. Integrity
 - C. Interoperability
 - D. Availability
12. An email filter is most effective against which types of attacks?
- A. Malicious code
 - B. Spamming
 - C. Spoofing
 - D. SYN floods
13. Prevention of fraud is embodied by all but which of the following activities?
- A. Job rotation
 - B. Mandatory vacations
 - C. Storage system quota management
 - D. Separation of duties
14. The * (star) Property is associated with which of the following security models?
- A. Clark-Wilson
 - B. Biba
 - C. Bell-LaPadula
 - D. Information Flow
15. Access control mechanisms operate by following which of these orders of security actions?
- A. Security policy, implementation, testing, and then tuning
 - B. Identification, authentication, authorization, and then accountability
 - C. Authentication, biometrics, token processing, and then auditing
 - D. Auditing, separation of duties, authorization, and then management
16. An authentication factor can be all but which of the following?
- A. Something you are
 - B. Something you have
 - C. Something you owe
 - D. Something you know
17. The simple integrity axiom of the Biba model can be simply stated by which of the following rules?
- A. No read down
 - B. No write down
 - C. No read up
 - D. No write up
18. What is a type II error of a biometric device?
- A. False rejection
 - B. False acceptance
 - C. Invalid enrollment
 - D. Interrupted authorization

628 PRACTICE EXAM

19. After a subject enters a pass phrase, what is created by the system and used to perform the actual authentication?
- A. One-time password
 - B. Virtual password
 - C. Single sign on password
 - D. Challenge token password
20. What is two-factor authentication?
- A. The process of typing in a username and a password
 - B. The use of a smart card
 - C. The use of two authentication factors
 - D. The use of a biometric device
21. Which of the following access control mechanisms is easiest to administer in an environment with a high personnel turnover rate?
- A. Access control lists
 - B. Rule-based access control
 - C. Role-based access control
 - D. Discretionary access control
22. Which of the following is the least secure?
- A. Challenge-response tokens
 - B. One-time passwords
 - C. Static passwords
 - D. Dynamic passwords
23. Accountability is provided through all but which of the following security mechanisms?
- A. Auditing
 - B. Lockout policy
 - C. Identification
 - D. Authentication
24. A user account name and an associated password are the most common representations of which of the following?
- A. Biometric enrollment
 - B. Identification and authentication
 - C. Two-factor authentication
 - D. Principle of least privilege
25. Kerberos is most effective against which of the following types of attack?
- A. Denial-of-service
 - B. Social engineering
 - C. Playback
 - D. Dictionary attacks
26. The most secure firewall is which of the following?
- A. Packet filtering firewall
 - B. Application gateway firewall
 - C. Kernel proxy firewall
 - D. Screened subnet firewall
27. An attack against wireless communications on a network involves violating which of the following?
- A. Confidentiality
 - B. Integrity
 - C. Availability
 - D. Throughput
28. SSL can be used to prevent which of the following types of attacks?
- A. Man-in-the-middle
 - B. Brute force and dictionary attacks

- C. Denial-of-service
D. Eavesdropping and hijacking
29. What is the most common reason a firewall has vulnerabilities?
- A. Use of multiple protocols
B. Use of discretionary access controls
C. Misconfiguration
D. Spoofed attacks waged against a network
30. Which type of firewall is easiest to implement?
- A. Static packet filter
B. Dynamic packet filter
C. Application gateway
D. Stateful inspection
31. PGP is a security mechanism that is effective against preventing which type of attack?
- A. Malicious code delivery
B. Denial-of-service
C. Email spoofing
D. Hijack attacks
32. VPNs with strong end-to-end encryption can be implemented using which of the following?
- A. Kerberos
B. SWIPE
C. PPTP
D. CHAP
33. Which of the following is considered a boundary security mechanism?
- A. Gateway
B. Firewall
C. Router
D. VPN
34. Which of the following forms of communication is essentially connectionless?
- A. Ethernet
B. TCP
C. Frame relay
D. ISDN
35. What is firewall security based on?
- A. Roles
B. Rules
C. Classifications
D. Sensitivity
36. Which of the following is a valid function for a firewall?
- A. Convert
B. Discard
C. Bounce
D. Broadcast
37. WAN connections, such as frame relay, ATM, and X.25, operate at which layer of the OSI model?
- A. Session
B. Network
C. Transport
D. Data Link
38. Token-Ring operates at which layer of the OSI model?
- A. Application
B. Session

630 PRACTICE EXAM

- C. Network
D. Physical
39. Routers operate at which layer of the OSI model?
- A. Application
B. Session
C. Network
D. Physical
40. Switches operate at which layer of the OSI model?
- A. Session
B. Network
C. Transport
D. Data Link
41. Routers provide a well-rounded security environment when used in combination with which of the following?
- A. Firewalls
B. Proxies
C. Gateways
D. Switches
42. Which of the following topologies can be used by both Ethernet and Token-Ring networks?
- A. Ring
B. Star
C. Bus
D. Mesh
43. The TCP/IP layer model has how many layers?
- A. 3
B. 4
C. 6
D. 7
44. Which networking topology generally requires the least amount of network cabling when connecting the same number of clients in a fixed pattern?
- A. Ring
B. Star
C. Bus
D. Mesh
45. Which of the following is true?
- A. UTP cabling includes a foil sheath.
B. EMI is reduced by increasing the twists per inch.
C. All twisted-pair wiring can be used up to 500 meters.
D. STP is impervious to tapping and eavesdropping.
46. All but which of the following are centralized remote access authentication systems?
- A. DIAMETER
B. TACACS+
C. RADIUS
D. CIRCUMFERENCE
47. What is the most common cause of network failures?
- A. Authentication database corruption
B. Network saturation
C. Denial-of-service attacks
D. Cabling problems

48. Sockets are associated with which of the following protocols?
- A. IGMP
 - B. TCP
 - C. IPX
 - D. SHTTP
49. What is another name for multi-port repeater?
- A. Switch
 - B. Router
 - C. Hub
 - D. Gateway
50. Which of the following cable types can be deployed in a single cable segment more than 200 meters in length?
- A. 10BASE-2
 - B. ThickNet
 - C. STP
 - D. 100BASE-T
51. What are the primary goals of security?
- A. Confidentiality, integration, and accessibility
 - B. Authentication, authorization, and accountability
 - C. Availability, integrity, and confidentiality
 - D. Physical, logical, and administrative
52. When evaluating risk, what is calculated by subtracting the applied countermeasures from the identified risks?
- A. Total risk
 - B. Residual risk
 - C. Controls gap
 - D. Acceptable risk
53. Which of the following is not a valid action that can be taken against risk when performing risk management?
- A. Reduce
 - B. Accept
 - C. Assign
 - D. Increase
54. What is acceptable risk?
- A. Cost of countermeasures > value of object
 - B. Cost of countermeasures < value of object
 - C. Attacker's cost > value of object
 - D. Attacker's cost < value of object
55. What is the process of deploying countermeasures to eliminate risk known as?
- A. Risk avoidance
 - B. Risk acceptance
 - C. Risk mitigation
 - D. Risk assignment
56. What is the level of risk an organization is willing to accept or assume to achieve a desired goal known as?
- A. Risk avoidance
 - B. Risk assignment
 - C. Risk mitigation
 - D. Risk tolerance
57. What is the proper definition of risk?
- A. Threat \times vulnerability
 - B. Threat \times controls gap
 - C. Vulnerability \times asset value
 - D. Vulnerability \times single loss expectancy

632 PRACTICE EXAM

58. Which of the following statements is not true?
- A. A purely quantitative risk analysis is possible.
 - B. A purely qualitative risk analysis is possible.
 - C. Quantitative assessment assigns real numbers to risks.
 - D. Qualitative assessment involves a fair amount of guesswork.
59. Which of the following is always an essential element of risk management?
- A. Deploying firewalls
 - B. Obtaining sign-off letters from management
 - C. Staying under budget
 - D. Applying desktop OS patches
60. Within an organization, which of the following offers optional instructions?
- A. Policies
 - B. Guidelines
 - C. Procedures
 - D. Standards
61. Which of the following is the most sensitive classification?
- A. Confidential
 - B. Top secret
 - C. Proprietary
 - D. Private
62. Which of the following are defined by entities outside the organization?
- A. Policies
 - B. Guidelines
 - C. Procedures
 - D. Standards
63. What does a trade secret do?
- A. Provides the owner exclusive rights for 17 years
 - B. Protects “original works of authorship”
 - C. Provides confidentiality of proprietary technical or business-related information
 - D. Establishes a word, name, symbol, color, sound, product shape, device, or combination of these used to identify and distinguish goods
64. When terminating an employee, which of the following is an important aspect of the removal process?
- A. Filing supply request forms
 - B. Reviewing nondisclosure agreements
 - C. Issuing the former employee new smart cards
 - D. Updating the former employee’s resume
65. In the CIA triad, availability means which of the following?
- A. Privacy
 - B. Timeliness
 - C. Consistency
 - D. Accuracy
66. What is another meaning for integrity?
- A. Privacy
 - B. Non-repudiation
 - C. Secret
 - D. Accessibility

67. Where does security management start?
- A. End users
 - B. System administrators
 - C. Company owner
 - D. Department manager
68. What is awareness a prerequisite of?
- A. Security certification
 - B. Security deployment
 - C. Security training
 - D. Security implementation
69. An organizational security policy should primarily focus on which activity?
- A. Hardware deployment
 - B. End user behavior modification
 - C. Software configuration
 - D. Data backups
70. To ensure proper coverage and application, an organization's security policies should be linked with which of the following?
- A. Countermeasures
 - B. Risks
 - C. Operating systems
 - D. User roles
71. When hiring new employees, what is an important part of educating them in regard to the organization's security policies and procedures?
- A. Training in a classroom environment
 - B. Posting the security policies on an intranet Web site
 - C. Having an executive teach the security awareness course
 - D. Obtaining a signed statement indicating they have read and understood the security policies and procedures
72. What is the primary reason organizational security policies are not followed?
- A. Difficult procedures
 - B. Adherence to strict public standards
 - C. Lack of enforcement
 - D. Cost of countermeasures
73. When defining security objectives, which of the following is the most important?
- A. The objective must be reasonable.
 - B. The objective must be achievable.
 - C. The objective must be effective.
 - D. The objective must be comprehensive.
74. Which of the following organizational security plans is usually useful, stable, and applicable for 1 year?
- A. Strategic plan
 - B. Operational plan
 - C. Tactical plan
 - D. Procedural plan
75. An operational plan can include all but which of the following?
- A. Project descriptions, including key milestones
 - B. The implementation schedule
 - C. Definitions of dependencies among strategies and a logical sequence of initiatives
 - D. Assessment of the current environment, such as risk assessment

634 PRACTICE EXAM

76. Which of the following is the most accurate description of a common computer virus?
- A. Malicious code that prevents legitimate activity from occurring on a system
 - B. Malicious code that replicates using a host program
 - C. An error on a hardware device that causes data corruption
 - D. An error caused by sending input to software of a volume larger than it was designed to handle
77. Privacy is easily compromised when which of the following is used on the Web?
- A. HTML
 - B. SSL
 - C. Cookies
 - D. Digital signatures
78. What are errors or problems encountered through the violation of data input block size known as?
- A. Buffer overflow
 - B. Flooding
 - C. Spoofing
 - D. Denial-of-service
79. Which of the following is not true in regard to software security?
- A. Security is often disabled for ease of installation.
 - B. Security must be configured for the specific environment.
 - C. Most software is secure right out of the box.
 - D. Modern software offers numerous features, and each must be evaluated in terms of security.
80. What can be the result of the failure of a programmer to properly handle software failures?
- A. System freezing or crashing (that is, a blue screen)
 - B. Resetting to default configuration
 - C. Elevation of auditing scope
 - D. Restarting into privileged mode
81. Database access is usually directed through a controlled client interface that provides which of the following?
- A. Availability and integrity
 - B. Confidentiality and integrity
 - C. Availability and authentication
 - D. Backups and redundancy
82. What is a mechanism that provides a structure for gathered data known as?
- A. A storage device
 - B. A database
 - C. A hierarchical relationship
 - D. A redundant array
83. What is a tuple?
- A. A table stored in a database
 - B. A row in a database
 - C. A collection of records of the same type
 - D. The attribute of one table that is the primary key of another table

84. What is the database component that holds the data that describes the database known as?
- A. A cell
 - B. The degree
 - C. The data dictionary
 - D. The schema
85. Which of the following statements is not true regarding a hierarchical data model?
- A. It combines records and fields that are related in a logical star structure.
 - B. Parents can have one child, many children, or no children.
 - C. It contains branches and leaves or data fields.
 - D. It's useful for mapping one-to-many relationships.
86. Which database model provides many-to-many relationships between elements?
- A. Relational data model
 - B. Hierarchical data model
 - C. Distributed data model
 - D. Inherent data model
87. A(n) _____ is an attribute in one relation that has values matching the primary key in another relation.
- A. Candidate key
 - B. Foreign key
 - C. Relation block
 - D. Element set
88. What is the cardinality of a database?
- A. The number of rows
 - B. The number of columns
 - C. The number of elements
 - D. The number of relationships
89. Within a database, a referential integrity mechanism is designed to perform which function?
- A. Upon an error, return the database to its previously saved state
 - B. Ensure that no record contains a reference to a primary key of a nonexistent record
 - C. Terminate a transaction and execute all changes made by an administrator
 - D. Verify that all structural and semantic rules are not violated
90. What is the ability of users to deduce information about data at higher sensitivity levels for which they do not have access privileges known as?
- A. Aggregation
 - B. Inference
 - C. Granularity
 - D. Escalation
91. What countermeasure can be used against the ability of users to deduce information about data at higher sensitivity levels for which they do not have access privileges?
- A. Database partitioning
 - B. Noise insertion
 - C. Polyinstantiation
 - D. Cell suppression
92. Which life cycle model allows for project modifications only to the preceding development stage within that cycle?
- A. Spiral model
 - B. Clark-Wilson model

636 PRACTICE EXAM

- C. Syngress model
D. Waterfall model
93. Which of the following should not be performed during the testing phase of a product development cycle?
- A. Test for handling of invalid input
B. Test using live or real field data
C. Test for handling of out-of-range values
D. Test using variations of conditions
94. Which level of the Software Engineering Institute's (SEI's) model for identifying the maturity of a software development process states that project practices are institutionalized?
- A. Level 1: Initiating
B. Level 2: Repeatable
C. Level 3: Defined
D. Level 4: Managed
95. A(n) _____ exhibits the reasoning capabilities similar to that of a human through the collection of rules and the building of inference mechanisms.
- A. Expert system
B. Computer program
C. Artificial intelligence
D. Neural network
96. Which of the following statements is true?
- A. An interpreted language is used to create precompiled applications.
B. Compiled code poses a higher security risk than interpreted code.
C. Java is a limited platform language.
D. An applet is a small program that is shared between numerous software packages simultaneously.
97. Restricting the flow of malicious code into your environment can take the form of all but which of the following?
- A. Screening applets and attachments at the firewall
B. Configuring Web browsers to refuse downloadable code
C. Accepting all digital certificates presented to your system
D. Training users about the threats of mobile code
98. A _____ is a type of malicious code that self-replicates to other systems and does not need a host program to function.
- A. Common virus
B. Worm
C. Trojan
D. Logic bomb
99. Which of the following is not considered a denial-of-service attack?
- A. Spoofing
B. Consuming bandwidth
C. Causing 100% CPU utilization
D. Redirecting legitimate traffic

100. Which of the following denial-of-service attacks takes the form of numerous incomplete initiations of the TCP three-way handshaking process?
- A. Smurf attack
 - B. Teardrop attack
 - C. Fraggle attack
 - D. SYN flood
101. Which of the following is not a goal of a cryptosystem?
- A. Confidentiality
 - B. Availability
 - C. Integrity
 - D. Non-repudiation
102. What is the data encryption standard (DES) an example of?
- A. An asymmetric key encryption algorithm
 - B. A symmetric key encryption algorithm
 - C. A non-repeating hash encryption algorithm
 - D. A repeating hash encryption algorithm
103. What is MD5 an example of?
- A. An asymmetric key encryption algorithm
 - B. A symmetric key encryption algorithm
 - C. A hash algorithm
 - D. A linear regression algorithm
104. IPSec provides protection of transmitted traffic using which two methods or modes?
- A. Linking and hashing
 - B. Transport and tunneling
 - C. Reporting and logging
 - D. Stateful and connectionless
105. The key length of _____ is 160 bits.
- A. MD5
 - B. SHA-1
 - C. MD2
 - D. 3DES
106. MD5 can be exploited using which type of attack?
- A. Dictionary
 - B. Scanning
 - C. Birthday
 - D. Spoofing
107. Tripwire is a well-known utility used for which purpose?
- A. Password database cracking
 - B. IDS
 - C. Manipulating ACLs
 - D. File integrity checking
108. The Public Key Infrastructure (PKI) is designed to provide or create a communications sharing environment that is which of the following?
- A. Restricted
 - B. Controlled
 - C. Trusted
 - D. Available
109. Proving the identity of both ends of a transaction using digital signatures, strong encryption algorithms, and the protection of private keys provides which of the following?
- A. Integrity
 - B. Trust

638 PRACTICE EXAM

- C. Confidentiality
D. Availability
110. Public Key Infrastructure (PKI) is most easily recognized as which of the following?
- A. A procedural guideline
B. A software product
C. An infrastructure
D. A hardware device
111. TLS and SSL can be used to protect all but one of the following types of traffic. Which one?
- A. FTP
B. Telnet
C. ICMP
D. Email
112. What does the Encapsulating Security Payload (ESP) component of IPSec provide?
- A. Non-repudiation
B. Limited authentication
C. Access control
D. Payload verification
113. Internet Key Exchange (IKE), which defines key management for IPSec, contains all but which of the following protocols?
- A. WTLS
B. ISAKMP
C. SKEME
D. Oakley Key Determination Protocol
114. Which of the following is a specific alternative to SSL for Web communications?
- A. SET
B. S-HTTP
C. PAP
D. S/MIME
115. Which of the following is a secure replacement for Telnet?
- A. S/MIME
B. TLS
C. SSH
D. WTP
116. Which protocol can be used to encrypt IEEE 802.11b communications?
- A. WEP
B. TLS
C. S/MIME
D. PKI
117. Which of the following is not an attack directed at cryptography?
- A. Brute force
B. Statistical
C. Birthday attack
D. Teardrop
118. Which of the following encryption algorithms is a replacement for DES?
- A. AES
B. SHA
C. MD5
D. RSA

119. DES, DSA, and ECDSA are all components of _____.
- A. DES
 - B. DSS
 - C. RSA
 - D. IPSec
120. Which of the following is a true statement about hashing algorithms?
- A. All use 128-bit hash values.
 - B. All are one-way functions.
 - C. All are very slow.
 - D. All process text in 1,024-bit blocks.
121. Which of the following can be used as a digital signature?
- A. DES
 - B. Blowfish
 - C. IDEA
 - D. El Gammal
122. Which of the following is a hashing algorithm?
- A. 3DES
 - B. Diffie-Hellman
 - C. HAVAL
 - D. ECC
123. When using a communications encryption system, what is the most important aspect of the cryptographic mechanism?
- A. Key strength
 - B. Useful lifetime
 - C. Key management
 - D. Cipher length
124. Which encryption scheme is unbreakable because each pass phrase or authentication code is used only once?
- A. Single sign on
 - B. One-way hash
 - C. Digital signatures
 - D. One-time pad
125. What is link encryption?
- A. An encryption system used to protect hyperlinks in a Web document
 - B. An encryption system that protects traffic only across a specific communications path
 - C. An encryption system that protects traffic from source to destination
 - D. An encryption system that protects traffic over VPNs
126. What is an outline of requirements necessary to properly support a specific security policy?
- A. A security model
 - B. A procedural manual
 - C. A proposal request
 - D. A standards document
127. Which of the following statements is not true?
- A. Security must be engineered.
 - B. Many aspects of the design and architecture of a system are dependent on security requirements.
 - C. Security should be added after the initial development of a system.
 - D. Security must be audited to be effective.

640 PRACTICE EXAM

128. A(n) _____ occurs if the operating system or the software fails to properly set boundaries and restrictions on how much data can be sent to the CPU.
- A. Denial of service
 - B. Buffer overflow
 - C. Data corruption
 - D. Encryption key disclosure
129. Nonvolatile storage (such as floppy disks, CD-ROM, and HDD) is labeled as which type of memory architecture?
- A. Primary storage
 - B. Secondary storage
 - C. Real storage
 - D. Virtual storage
130. What type of memory is also known as firmware?
- A. BIOS
 - B. RAM
 - C. ROM
 - D. EPROM
131. What is the most trusted physical component of a computer?
- A. RAM
 - B. Storage devices
 - C. Motherboard/mainboard
 - D. CPU
132. Software uses virtual memory managed by a memory mapper component (that is, virtual memory manager) in the kernel. Why is this done?
- A. It provides for faster memory usage.
 - B. It reduces system overhead.
 - C. Software is not trusted.
 - D. Hardware can't directly support sufficient physical RAM for most software products.
133. Which ring of the protection ring model is designated for input and output device drivers?
- A. Ring 0
 - B. Ring 1
 - C. Ring 2
 - D. Ring 3
134. Which of the following statements about the protection ring model is not true?
- A. If an entity needs to access a resource in a ring of greater protection, a system call is executed.
 - B. The higher the number, the greater the protection provided within that ring.
 - C. Entities can access resources only within their ring and in rings of lower protection.
 - D. Rings are used to designate protection levels for various aspects of the software components (kernel, drivers, utilities, application, and so on) of a computer.
135. The operating state labeled "problem state" is identified as which of the following conditions?
- A. An application is executing.
 - B. An application is ready to resume execution.
 - C. A system level or privileged operation is underway.
 - D. An error has occurred.
136. What is multitasking?
- A. Opening several applications at once
 - B. Processing more than one thread at once

- C. Processing more than one process at once
D. Using more than one processor to execute instructions in parallel
137. Which of the following is true?
- A. The more complex a security system is, the less assurance it provides.
 - B. Protection must occur at the data end of a resource request.
 - C. No security measure can regulate activities between programs and objects.
 - D. The simpler a security system is, the less security it can provide.
138. What is the Trusted Computing Base (TCB)?
- A. A fully secured computer system from a vendor
 - B. The collection of components within a system that provides a specific level of trust (that is, security)
 - C. The hardware components of a computer
 - D. The software components used to implement the security policy
139. To enforce accountability, a system must provide which of the following?
- A. Hardware segmentation
 - B. Resource isolation
 - C. Inference prevention
 - D. Tuple exploitation
140. The security model represented by a directed graph that specifies the rights a subject can transfer to an object or that a subject can obtain from another subject is known as which of the following?
- A. Information flow model
 - B. Take-Grant model
 - C. Clark-Wilson model
 - D. Inheritance model
141. A state machine can be labeled as such if all but which of the following is true?
- A. Always boots into a secure state
 - B. Executes commands securely
 - C. Allows for a wide variation of transactions
 - D. Restricts the subject to accessing objects only by means that are prescribed by the security policy
142. What are the rows of an access matrix known as?
- A. Access control lists
 - B. Inheritance lists
 - C. Capability lists
 - D. Authorization lists
143. Which of the following is not a weakness of the Bell-LaPadula model?
- A. Does not consider covert channels
 - B. Does not consider network-based resource/object sharing
 - C. Does not explicitly define what a secure state transaction actually means
 - D. Does not protect the confidentiality of data
144. Which of the following models is lattice based?
- A. Biba model
 - B. Clark-Wilson model
 - C. Take-Grant model
 - D. Information Flow model

642 PRACTICE EXAM

145. The Clark-Wilson model is primarily concerned with which of the following?
- A. Prevention of unauthorized disclosure of data
 - B. Prevention of unauthorized modification of data
 - C. Prevention of inability to access data in a timely fashion
 - D. Prevention of data inference
146. Separation of duties is a foundational element of which security model?
- A. Biba model
 - B. Clark-Wilson model
 - C. Bell-LaPadula model
 - D. Information Flow model
147. The Trusted Computer System Evaluation Criteria (TCSEC) is defined in which publication?
- A. Red Book
 - B. Purple Book
 - C. Yellow Book
 - D. Orange Book
148. Which of the following is a replacement and an update to Trusted Computer System Evaluation Criteria (TCSEC)?
- A. Trusted Database Management System (TDI)
 - B. Common Criteria (CC)
 - C. Trusted Network Interpretation (TNI)
 - D. Information Technology Security Evaluation Criteria (ITSEC)
149. According to Trusted Computer System Evaluation Criteria (TCSEC), which of the following is the highest security valuation?
- A. A
 - B. B
 - C. C
 - D. D
150. Which TCSEC security label requires the use of security domains?
- A. C1
 - B. B3
 - C. A1
 - D. D
151. Which TCSEC security designation is the highest possible that still allows for the presence of covert channels?
- A. C2
 - B. B1
 - C. B2
 - D. A1
152. Which National Information Assurance Certification and Accreditation Process (NIACAP) accreditation type is used to evaluate a specific self-contained location?
- A. Type
 - B. Site
 - C. Domain
 - D. System

153. A closed system architecture has all features or characteristics except for which of the following?
- A. Published specifications
 - B. Proprietary
 - C. Offers security through obscurity
 - D. No significant third-party support
154. Which of the following is a type of covert channel?
- A. Side band modem line
 - B. Timing
 - C. Encrypted removable media
 - D. PGP protected email
155. Which of the following is not a valid counter-measure for preventing the use of a backdoor?
- A. Network-based IDS
 - B. Use of strict file system access controls
 - C. Use of communication encryption protocols
 - D. Auditing system activities
156. What is the security condition in which no single person has complete access to or control over all the security mechanisms on a system known as?
- A. Preventative control
 - B. Separation of duties
 - C. Detective control
 - D. Access control
157. Which of the following reduces the probability of collusion between employees to perform fraudulent activities?
- A. Separation of duties
 - B. Detective controls
 - C. Rotation of duties
 - D. Two-man controls
158. What is the primary goal of security configuration management?
- A. To ensure that all changes made to a system do not result in reduced security
 - B. To ensure that changes made to a system are performed only by authorized administrators
 - C. To track the activities of administrators' use of elevated privileges
 - D. To prevent end users from performing administrative tasks
159. Change management should provide for all but which of the following?
- A. Tracking and approving all changes to a system
 - B. Reducing negative effects on productive use of the system
 - C. Documenting changes to system security
 - D. Preventing rollback to a previous version of the system
160. Which of the following is not an appropriate change management procedure?
- A. Catalog the intended change.
 - B. Schedule the change.
 - C. Evaluate the change in light of industry security standards.
 - D. Report the change appropriately.
161. Which of the following is not a valid procedure for managing personnel security?
- A. Skills assessment exams
 - B. Background checks
 - C. Mandatory one-week vacation increments
 - D. Separation of duties

644 PRACTICE EXAM

162. An owner of an organization will be held liable for costs associated with a security breach causing a loss if he is unable to _____.
- A. Produce a security policy
 - B. Show due care
 - C. Identify a firewall deployment
 - D. Reference a list of job responsibility designations
163. What is piggybacking?
- A. When a person walks through a secured doorway without self-authenticating immediately behind someone who performed proper self-authentication
 - B. Replaying the packets of a captured session to restart the communication process
 - C. Adding malicious code to an email or a document
 - D. Connecting to an open port over a VPN connection
164. What is the data that is still present on a storage device after it has been erased known as?
- A. Bad sectors
 - B. Recycled contents
 - C. Data remnants
 - D. File allocation table residue
165. Security controls should be which of the following?
- A. As complex as possible
 - B. As exhaustive as possible
 - C. As transparent to the user as possible
 - D. As restrictive as possible
166. The goals of monitoring and auditing are all but which of the following?
- A. Resolution of problems
 - B. Identification of abnormalities
 - C. Prevention of attacks
 - D. Identification of normal events
167. What is the monitoring activity that obtains information simply by asking for it known as?
- A. Sniffing
 - B. Dumpster diving
 - C. Social engineering
 - D. Demon dialing
168. What is a clipping level?
- A. The point at which too much data is gathered by an auditing system and events are lost.
 - B. The level below which all normal activities occur. Only events above this level should be suspect.
 - C. The level at which too much data is being transmitted over a network (that is, complete saturation and consumption of available bandwidth) and traffic is lost.
 - D. The point at which an intruder in a honeypot or padded cell is automatically disconnected.
169. The use of a clipping level allows for all but which of the following activities?
- A. Detection of slow, low-profile intrusion attempts against a system
 - B. Detection of high-occurrence repetitive mistakes by a user
 - C. Detection of users who are attempting to exceed their authorization levels
 - D. Detection of high-traffic directed intrusion attempts

170. Which of the following is not true?
- A. Audit logs should be retained for historical reference.
 - B. Audit logs should be protected from alteration.
 - C. Audit logs should be capable of recording data during an event (in other words, 100% availability).
 - D. Audit logs should be stored only on removable media.
171. Which of the following is not a threat from inappropriate activities?
- A. End users accessing pornographic, political, religious, or violent content
 - B. Managers conducting private business
 - C. System operators discussing confidential material with non-employees
 - D. Program designers including omission errors in their custom scripts
172. All but which of the following are valid countermeasures to traffic analysis vulnerabilities?
- A. Use of encryption
 - B. Message padding
 - C. Noise transmission
 - D. Analysis of covert channels
173. Which of the following is a vulnerability scanning tool?
- A. TCPwrappers
 - B. SATAN
 - C. nmap
 - D. Back Orifice
174. Countermeasures for port mapping attacks include all but which of the following?
- A. Filtering traffic at a firewall
 - B. Disabling banners on network services
 - C. Deploying a strong password policy
 - D. Deploying an IDS
175. A _____ program is designed to recover from a system freeze or malfunction by bypassing security and access controls.
- A. Smurf
 - B. Superzapping
 - C. Sniffer
 - D. SATAN
176. Sniffers that support decoding capabilities are able to perform which activity?
- A. Detect intrusion attempts
 - B. Store their capture buffers on a storage device
 - C. Reveal the contents of captured traffic
 - D. Edit packets and retransmit them
177. Which of the following is not a sniffer utility?
- A. John the Ripper
 - B. Snort
 - C. Trinix
 - D. nmap
178. What is a countermeasure for session hijacking performed using spoofed IP addresses or the Juggernaut or Hunt utility?
- A. Two-factor authentication
 - B. Role-based access controls
 - C. Event auditing
 - D. Use IPSec authentication

646 PRACTICE EXAM

179. Which of the following terms is used to label or describe a minor disruptive event where an organization must recover and continue to support critical functions?
- A. Disaster recovery planning
 - B. Business continuity planning
 - C. Backup restoration planning
 - D. Security policy planning
180. Which of the following is not a factor of business continuity planning?
- A. Provides a means to upgrade security mechanisms
 - B. Reduces the risk of financial loss
 - C. Mitigates risks associated with the disruptive event
 - D. Recovers from problems quickly
181. When a disaster occurs, which of the following is the most important and primary activity that should occur?
- A. Locate the off-site backup copies.
 - B. Order replacement hardware.
 - C. Ensure that all personnel are accounted for.
 - D. Issue a press release regarding the disaster.
182. Who is ultimately responsible for the success of a business continuity plan?
- A. Security administrators
 - B. End users
 - C. Deployment operatives
 - D. Senior management
183. Which of the following is not one of the three primary goals of business impact analysis (BIA)?
- A. Downtime estimation
 - B. Criticality prioritization
 - C. Vulnerability assessment
 - D. Resource requirements
184. Which of the following is a key element in the implementation process of a business continuity plan?
- A. Industry standards
 - B. Employee awareness
 - C. Dry run testing
 - D. Senior management approval
185. Which of the following should be true of an organization's business continuity plan?
- A. There should be only one.
 - B. Once developed, the plan requires no maintenance.
 - C. Auditing the plan is unnecessary.
 - D. Each department should have its own local plan.
186. Disaster recover planning should address all but which of the following?
- A. Paying investors recovery dividends
 - B. Providing backup operations during the recovery process
 - C. Providing for a salvage operation after the primary recovery is complete
 - D. The procedures necessary to respond to an emergency

187. Which type of subscription service site offers a computer facility readily available with electricity, air conditioning, and computers but does not have applications installed?
- A. Hot site
 - B. Warm site
 - C. Cold site
 - D. Secondary site
188. Which of the following is not a disadvantage of a hot site?
- A. Low administration overhead.
 - B. Expense.
 - C. Service providers often oversell their capabilities.
 - D. Contains a real-time mirrored image of production data.
189. Which of the following is not considered an adequate resource for disaster recovery?
- A. Hot site
 - B. Warm site
 - C. Cold site
 - D. Secondary site
190. When selecting an offsite facility for use during disaster recovery, which of the following is the most important aspect to consider?
- A. Cost
 - B. Square footage
 - C. Distance from original site
 - D. Exclusive use
191. The process of backing up data to an offsite location is known as which of the following?
- A. Remote storage
 - B. Electronic vaulting
 - C. Warm site development
 - D. Database shadowing
192. What is remote journaling?
- A. Duplicating data sets to multiple servers
 - B. Batch processing transactions to an alternative site
 - C. Parallel processing of transactions to an alternative site
 - D. Transmitting data to an alternative site via WAN connections
193. Which of the following is true of disaster recovery plans?
- A. Testing can be performed by any means.
 - B. Demonstrated recovery capability exists even without testing.
 - C. Tests only need to involve critical components of the plan.
 - D. If a plan is not tested, it does not work.
194. Which type of test involves the distribution of the plan to all appropriate personnel for review?
- A. Checklist test
 - B. Structured walk-through test
 - C. Simulation test
 - D. Parallel test

648 PRACTICE EXAM

195. Which type of test is a full test but the activities at the production environment are not stopped?
- A. Structured walk-through test
 - B. Simulation test
 - C. Parallel test
 - D. Full-interruption test
196. Which type of test works through the recovery plan up to the point just before alternative processing is initiated?
- A. Checklist test
 - B. Structured walk-through test
 - C. Simulation test
 - D. Parallel test
197. The _____ team returns to the original site only after the possibility of personal injury is eliminated.
- A. Recovery
 - B. Salvage
 - C. Response
 - D. Evaluation
198. When is an emergency actually over?
- A. When personal danger is eliminated
 - B. When operations are fully functional at an alternative site
 - C. When the organization fully returns to the original site
 - D. When all critical functions are supported
199. When recovering from a disaster, what should be performed first?
- A. Restore the least critical functions.
 - B. Restore critical functions.
 - C. Salvage equipment from the original site.
 - D. Evaluate public relations damage.
200. What is an important item that should be part of a disaster recovery plan but is often overlooked?
- A. Designation of an alternative site in the event the primary site is destroyed
 - B. Adequate backup of data
 - C. Quick restoration of business processes
 - D. Continuing to pay employees even if business production is interrupted
201. Which of the following is not restricted in the (ISC)² Code of Ethics?
- A. Acting dishonestly
 - B. Writing viruses
 - C. Providing incompetent service
 - D. Detracting from the security profession
202. Which of the following is not considered an unethical activity by the Internet Activities Board (IAB) according to RFC 1087?
- A. Gaining unauthorized access to resources on the Internet
 - B. Wasting resources
 - C. Selling products over the Internet
 - D. Compromising the privacy of users
203. Which of the following is not part of the Generally Accepted Systems Security Principles (GASSP)?
- A. The mission of an organization should be supported by the security policy.
 - B. Sound management has a foundation of security principles.

- C. Computer security should be cost effective.
D. System security can't be bound by societal restraints or factors.
204. Which of the following is not considered a computer crime?
A. Wasting resources
B. Password theft
C. Emanation eavesdropping
D. Distribution of malicious code
205. TEMPEST is used for what purposes?
A. Reading all email transmitted over the Internet
B. Retaining a copy of every Web site on the Internet
C. Preventing the interception of RF emanations
D. Tracking messages on the Internet for key phrases
206. What is pretending to be someone else to gain a greater level of access known as?
A. Espionage
B. Masquerading
C. Scripting
D. Superzapping
207. The theft of small amounts of information from numerous sources to reveal or extract highly confidential information is known as which type of attack?
A. Salami
B. Birthday
C. Sniffing
D. Spoofing
208. Modifying data through unauthorized means is known as which of the following?
A. Masquerading
B. Social engineering
C. Data diddling
D. Superzapping
209. Which of the following is not a significant restriction to the investigation of computer crimes?
A. Intangibility of evidence.
B. Evidence gathering requires no special skills.
C. Compressed investigational time frame.
D. Investigations might interfere with normal system operations and productivity.
210. In 1991, the U.S. Federal Sentencing Guidelines were revised in regard to punishments for breaking federal laws so that the severity of punishment is a direct relation to the degree _____.
A. The organization demonstrates due diligence
B. The perpetrator demonstrates technical expertise
C. Of loss of public confidence and profitability
D. Of the actual damage incurred
211. Which of the following is not an important aspect of showing due care?
A. Creating disaster recovery and business continuity plans
B. Implementing data backups and providing for hardware replacement
C. Public access to periodic vulnerability assessments
D. Intelligent use of physical and logical access controls

650 PRACTICE EXAM

212. A legal liability for the implementation of a safeguard or countermeasure is demonstrated based on which of the following?
- A. If the cost of the countermeasure is more than the cost of the boundary protection mechanisms
 - B. If the cost of the vulnerability is less than the cost of the safeguard
 - C. If the estimated cost of out-of-court settlements is more than the cost of the safeguard
 - D. If the cost of the countermeasure is less than the expected loss from an exploited vulnerability
213. What is the rule of the 1991 U.S. Federal Sentencing Guidelines that states that senior officials must perform their duties with the same care that ordinary sensible people would exercise under similar circumstances known as?
- A. The due care rule
 - B. The accountability rule
 - C. The golden rule
 - D. The prudent man rule
214. For negligence on the part of senior executives in the event of a disaster to be proven, which of the following must be demonstrated?
- A. Insufficient due diligence
 - B. A legally recognized obligation
 - C. Lack of applicable industry standards
 - D. No personnel injury occurred
215. Which of the following bodies of law are based on precedent?
- A. Statutory law
 - B. Administrative law
 - C. Civil law
 - D. Common law
216. Which body of law is directed toward the protection of the public and can offer punishments of financial penalties and imprisonment?
- A. Civil law
 - B. Criminal law
 - C. Regulatory law
 - D. Statutory law
217. Which element of intellectual property law provides the creator of a work exclusive rights for 17 years?
- A. Patent
 - B. Copyright
 - C. Trade secret
 - D. Trademark
218. European privacy laws are _____ U.S. privacy laws.
- A. Less restrictive than
 - B. More restrictive than
 - C. About the same as
 - D. Based on
219. When implementing electronic monitoring of all email on a company network, all but which of the following must be true?
- A. Monitoring is applied equally to all persons.
 - B. All users are informed of the network's acceptable use policy.
 - C. Details about who will read email and how long email will be backed up must be provided.
 - D. Users are provided a guarantee of privacy.

220. The act of encouraging the commission of a crime by an individual who initially had no intention of committing a crime is known as which of the following?
- A. Entrapment
 - B. Enticement
 - C. Entertainment
 - D. Espionage
221. A computer incident response team (CIRT) is responsible for all but which of the following?
- A. Reducing risk after an incident
 - B. Gathering evidence related to an incident
 - C. Minimizing negative impact on public relations due to an incident
 - D. Purging audit logs of details related to an incident
222. Which of the following should be performed during the initial process of evidence gathering at the scene of a computer crime?
- A. Reboot the system
 - B. Image the hard drive
 - C. Turn off power supplies
 - D. Use a portable x-ray device to scan the contents of the computer boxes
223. Evidence must be all but which of the following?
- A. Relevant
 - B. Permissible
 - C. Sufficient
 - D. Reliable
224. When gathering evidence of a computer crime, printouts should be identified or labeled using what?
- A. Removable stickers
 - B. Permanent markers
 - C. Pencils
 - D. Hole punches
225. What is the most important aspect of evidence gathering?
- A. Proper labeling
 - B. Prevention of alteration or tampering
 - C. Return of evidence to owner
 - D. Enclosure in an air-tight container
226. What type of evidence proves or disproves a specific act through oral testimony based on evidence gathered through the witness's five senses?
- A. Direct evidence
 - B. Best evidence
 - C. Circumstantial evidence
 - D. Hearsay evidence
227. What is evidence that is not based on personal, firsthand knowledge of the witness but was obtained from another source known as?
- A. Circumstantial evidence
 - B. Opinions
 - C. Hearsay evidence
 - D. Secondary evidence

652 PRACTICE EXAM

228. Which type of evidence is generally inadmissible in court?
- A. Hearsay evidence
 - B. Direct evidence
 - C. Circumstantial evidence
 - D. Secondary evidence?
229. Gathering or discovering enough evidence about a subject to consider an individual a suspect is known as which of the following?
- A. Conducting an interview
 - B. Conducting an audit
 - C. Conducting an interrogation
 - D. Conducting an assessment
230. Which of the following applies to federal agencies and is directed toward the protection of information about private individuals that is stored in government databases?
- A. Paperwork Reduction Act of 1995
 - B. U.S. Computer Fraud and Abuse Act
 - C. Gramm-Leach-Bliley Act of 1999
 - D. U.S. Privacy Act of 1974
231. The act of training users about an organization's security policy is which type of control?
- A. Physical
 - B. Logical
 - C. Administrative
 - D. Technical
232. Fire detection and suppression operations are which type of security control?
- A. Physical
 - B. Logical
 - C. Administrative
 - D. Technical
233. Which of the following is not an important security concern to evaluate when selecting a new site location?
- A. Local crime rate
 - B. Property tax rate
 - C. Police, medical, and fire services
 - D. Hazards from the surrounding area
234. When evaluating the safety and security of a facility, which of the following is not an important consideration?
- A. Combustibility
 - B. Crawl space
 - C. Load rating
 - D. Proximity to telephone company
235. What are the benefits of using human-incompatible server and equipment areas?
- A. Better fire suppression systems
 - B. Improved temperature controls
 - C. Optimized use of space
 - D. More efficient emergency protection area for personnel
236. Which of the following physical security threats are not violations of availability?
- A. Computer service interruptions
 - B. Unauthorized disclosure
 - C. Physical damage to hardware
 - D. Theft of equipment

237. Security controls must always do which of the following?
- A. Provide an impenetrable border
 - B. Be invisible to the user
 - C. Comply with laws and regulations
 - D. Protect data accessibility
238. Maintaining system availability is advanced by replacing hardware when which of the following occurs?
- A. As it reaches its mean time to repair
 - B. As it reaches its mean time between failures
 - C. As it reaches a six-month active service life-time
 - D. As the budget allows
239. What is momentary low voltage known as?
- A. Fault
 - B. Sag
 - C. Brownout
 - D. Noise
240. Traverse mode noise is the EMI generated by which of the following?
- A. The difference between hot and neutral wires
 - B. The difference between ground and neutral wires
 - C. The difference between hot and ground wires
 - D. The difference between hot wires of different devices
241. What is the ideal operating humidity for computer components?
- A. 20%–40%
 - B. 40%–60%
 - C. 60%–80%
 - D. 80%–100%
242. A static electricity voltage of what level will cause a system shutdown?
- A. 40
 - B. 1,500
 - C. 2,000
 - D. 17,000
243. Which type of fire extinguisher should be used for electrical fires?
- A. Class A
 - B. Class B
 - C. Class C
 - D. Class AB
244. Which of the following types of sprinkler systems is most recommended for computer centers?
- A. Dry pipe
 - B. Wet pipe
 - C. Deluge pipe
 - D. Preaction pipe
245. Which gas is primarily used to replace Halon in fire suppression systems?
- A. FM-100
 - B. FM-200
 - C. Halix
 - D. CO₂

654 PRACTICE EXAM

246. The benefits of guards for maintaining a physical security perimeter include all but which of the following?
- A. Ability to adjust to quickly changing conditions
 - B. Available for a nearly infinite variety of environments and conditions
 - C. Able to recognize intrusion patterns in real time
 - D. Able to make value judgments based on subjective information
247. Dogs are often a more suitable alternative to guards for numerous reasons, such as?
- A. Cost
 - B. Reliability
 - C. Maintenance
 - D. Liability issues
248. What is a mantrap?
- A. A double set of doors often monitored by a guard
 - B. A type of encryption algorithm
 - C. A fence surrounding a secure facility
 - D. A perimeter traffic monitor
249. Which is the most common form of perimeter or boundary protection?
- A. Dogs
 - B. Guards
 - C. CCTV
 - D. Lighting
250. What is the act of degaussing and overwriting data media for intended use outside the protected and secured environment known as?
- A. Destruction
 - B. Purging
 - C. Cleaning
 - D. Data mining

Answers to Exam Questions

1. **B.** The principle of least privilege implies users are granted minimal necessary access to perform their work tasks.
2. **C.** Mandatory access control must have subject classification to control access. Discretionary, ACLs, and rule-based all employ object-specific controls.
3. **C.** Spoofing is an active attack.
4. **B.** Administrative, logical, and physical controls are mechanisms of access control.
5. **B.** Padded cells include a simulated environment, logging capabilities, and malicious action restrictions, but they do not contain confidential data.
6. **C.** IDS is a detective security measure; it looks for abnormal or unauthorized activity. IDS does not prevent attacks directly, but it does inform system administrators of weaknesses that should be patched. IDS is usually not reactive or corrective. Some newer IDS products offer moderate reactive activities, such as disabling breached ports, but the CISSP CBK still defines IDS as detective only.
7. **B.** IDS is weakest at detecting spoofing attacks.

8. **C.** The crossover error rate (CER) is the performance rating for biometric devices that is used to judge the relative effectiveness between similar devices from different vendors.
9. **B.** Legal ramifications are the most important aspect to consider when deploying a honeypot.
10. **A.** The Bell-LaPadula security model was designed to address confidentiality.
11. **B.** The Biba security model was designed to address integrity.
12. **B.** Email filters are most effective against spamming attacks.
13. **C.** Storage system quota management is not a form of fraud prevention. Job rotation, mandatory vacations, and separation of duties are all forms of fraud prevention.
14. **C.** The * (star) property is associated with the Bell-LaPadula security model.
15. **B.** The order of security actions performed by access control mechanisms is identification, authentication, authorization, and then accountability.
16. **C.** Something you owe is not a valid authentication factor.
17. **A.** The simple integrity axiom can be simply stated as no read down.
18. **B.** A type II error is a false acceptance.
19. **B.** A virtual password is created from a pass phrase that is used for the actual authentication process.
20. **C.** Two-factor authentication is the use of any two authentication factors.
21. **C.** Role-based access control is the easiest to administer for environments with high personnel turnover rates. Role-based access control assigns privileges to roles instead of individuals. In environments with a high rate of turn over, assigning roles to new users is easier than modifying ACLs (which are discretionary controls) or altering rules.
22. **C.** Static passwords are the least secure password mechanism.
23. **B.** Lockout policy does not provide accountability.
24. **B.** A username and password are the most common representations of identification and authentication.
25. **C.** Kerberos is most effective against playback attacks.
26. **D.** A screened subnet firewall is the most secure because it employs a screened subnet within which the bastion host firewall resides. This effectively adds another layer of protection the other three firewall types do not offer.
27. **A.** Confidentiality is primarily violated when an attack is waged against wireless communications.
28. **D.** SSL can be used to prevent eavesdropping and hijacking attacks.
29. **C.** A firewall's vulnerabilities are most often caused by misconfiguration.
30. **A.** A static packet filter firewall is the easiest to implement.
31. **C.** PGP is effective against preventing email spoofing attacks.
32. **C.** PPTP can be used to implement a VPN with strong end-to-end encryption.

656 PRACTICE EXAM

33. **B.** A firewall is a boundary security mechanism.
34. **A.** Ethernet is a connectionless communication form. TCP, frame relay, and ISDN are all connection-oriented communication forms.
35. **B.** Firewall security is based on rules.
36. **B.** Discard is a valid function of a firewall.
37. **B.** WAN connections operate at the Network layer (layer 3).
38. **D.** Token-Ring operates at the Physical layer (layer 1).
39. **C.** Routers operate at the Network layer (layer 3).
40. **D.** Switches operate at the Data Link layer (layer 2).
41. **A.** Firewalls and routers provide a well-rounded security environment when used together.
42. **B.** A star topology can be used by both Ethernet and Token-Ring networks.
43. **B.** The TCP/IP layer model has 4 layers.
44. **B.** A star topology generally requires the least amount of network cabling.
45. **B.** EMI is reduced by increasing the twists per inch.
46. **D.** CIRCUMFERENCE is not a centralized remote access authentication system.
47. **D.** Cabling problems are the most common cause of network failures.
48. **B.** Sockets or ports are associated with TCP.
49. **C.** A hub is a multi-port repeater.
50. **B.** ThickNet, or 10BASE-5, can be deployed 500 meters.
51. **C.** The primary goals of security as defined by the CIA Triad are availability, integrity, and confidentiality.
52. **B.** Identified risk minus countermeasures is residual risk.
53. **D.** Increasing risk is not a valid action within risk management.
54. **A.** Acceptable risk occurs when the cost of countermeasures exceeds the value of the object.
55. **C.** Risk mitigation is the process of deploying countermeasures.
56. **D.** Risk tolerance is the level of risk an organization is willing to accept or assume to achieve a desired goal.
57. **A.** Risk can be defined as threat \times vulnerability. The control's gap is the benefit gained by implementing safeguards. It is the reduction of risk—it is not used to calculate risk. Risk is also not a product of an asset value or SLE.
58. **A.** A purely quantitative risk analysis is not possible because you can't quantify a qualitative item.
59. **B.** Obtaining sign-off letters from management is always an essential element of risk management.
60. **B.** Guidelines provide optional instructions within an organization.
61. **B.** Top secret is the most sensitive classification.
62. **D.** Standards are defined by entities outside the organization.
63. **C.** A trade secret provides confidentiality of proprietary technical or business-related information.
64. **B.** An important aspect of the removal process is to remind the former employee about your non-disclosure agreements.
65. **B.** Availability, within the CIA triad, can also mean timeliness.
66. **B.** Integrity can also mean non-repudiation.

67. **C.** Security management starts with the company owner.
68. **C.** Awareness is a prerequisite of security training.
69. **B.** A security policy should primarily focus on end user behavior modification.
70. **B.** An organization's security policies should be linked to risks.
71. **D.** An important part of new employee education is to obtain a signed statement indicating the employee has read and understood the security policies and procedures.
72. **C.** Lack of enforcement is the primary factor why organizational security policies are not followed.
73. **B.** The most important aspect of defining security objectives is that the object must be achievable.
74. **C.** The tactical plan is usually useful, stable, and applicable for only about 1 year.
75. **D.** An operational plan does not include assessment of the current environment, such as risk assessment.
76. **B.** A common virus is malicious code that replicates using a host program.
77. **C.** The use of cookies often compromises privacy.
78. **A.** Violation of data input block size is a buffer overflow.
79. **C.** Software is rarely secure right out of the box.
80. **D.** Restarting into privileged mode is a possible result if software failures are not properly managed by program developers.
81. **B.** Database access is usually directed through a controlled client interface that provides confidentiality and integrity.
82. **B.** A mechanism that provides structure for gathered data is known as a database.
83. **B.** A tuple is a row in a database.
84. **D.** The schema is the database component that holds the data that describes the database.
85. **A.** A hierarchical data model combines records and fields that are related in a logical tree structure, not a star.
86. **C.** A distributed data model provides for many-to-many relationships between elements.
87. **B.** A foreign key is an attribute in one relation that has values matching the primary key in another relation.
88. **A.** The cardinality of a database is the number of rows.
89. **B.** A referential integrity mechanism is designed to ensure that no record contains a reference to a primary key of a nonexistent record.
90. **B.** The ability of users to deduce information about data at higher sensitivity levels for which they do not have access privileges is known as inference.
91. **A.** Database partitioning is the countermeasure to prevent inference.
92. **D.** The waterfall model allows for project modifications only to the preceding development stage.
93. **B.** Live or real field data should never be used to test products.
94. **B.** Level 2, the Repeatable level, of the SEI project process maturity scale states that project practices are institutionalized.
95. **A.** An expert system exhibits the reasoning capabilities similar to that of a human through the collection of rules and the building of inference mechanisms.

658 PRACTICE EXAM

96. **B.** Compiled code poses a higher security risk than interpreted code because malicious code can be embedded in the compiled code and be difficult to detect.
97. **C.** Accepting all digital certificates presented to your system is not a mechanism for restricting malicious code. Digital signatures can be falsified or have untrusted backing and thus provide an unrestricted path into your system.
98. **B.** A worm is a type of malicious code that self-replicates to other systems and does not need a host program to function.
99. **A.** Spoofing is not considered a denial-of-service attack; it is an attack type of its own. Spoofing is the impersonation of something other than who you are.
100. **D.** A SYN flood is a denial-of-service attack that takes the form of numerous incomplete initiations of the TCP three-way handshaking process.
101. **B.** Availability is not a goal of cryptosystems; authenticity is.
102. **B.** DES is an example of a symmetric key encryption algorithm.
103. **C.** MD5 is an example of a hash algorithm.
104. **B.** IPsec uses the transport and tunneling modes.
105. **B.** SHA-1 has a key length of 160 bits.
106. **C.** MD5 can be exploited using the birthday attack.
107. **D.** Tripwire is a file integrity checking utility.
108. **C.** The goal of PKI is to create trusted environments.
109. **B.** Proving identities provides trust.
110. **C.** PKI is an infrastructure.
111. **C.** ICMP can't be protected by TLS or SSL.
112. **B.** ESP provides limited authentication.
113. **A.** WTLS is a wireless encryption protocol, not part of IPsec's IKE.
114. **B.** S-HTTP is an alternative to SSL. S-HTTP offers Web communication protection by encrypting individual documents rather than the entire session.
115. **C.** SSH, or Secure Shell, is a secure replacement for Telnet.
116. **A.** WEP, or Wired Equivalent Privacy protocol, is used to encrypt IEEE 802.11b (wireless) communications.
117. **D.** Teardrop is a DoS attack and is not aimed at cryptography.
118. **A.** AES is a replacement for DES. DES is an older standard based on 56-bit keys and is easily broken by current technology. AES is a very strong and very fast replacement. AES is based on the Rijndael algorithm and uses 128-, 192-, or 256-bit keys.
119. **B.** DES, DSA, and ECDSA are all components of the Digital Signature Standard (DSS).
120. **B.** All hash algorithms are one-way functions.
121. **D.** El Gamal, an asymmetric key algorithm, can be used as a digital signature.
122. **C.** HAVAL is a hashing algorithm.
123. **C.** Key management is the most important aspect of a cryptographic system. Without proper key management, none of the other elements of an encryption communication system matter.
124. **D.** A one-time pad is the encryption scheme that is unbreakable because each pass phrase or authentication code is used only once.

125. **B.** Link encryption is an encryption system that protects traffic only across a specific communications path.
126. **A.** A security model is an outline of requirements necessary to properly support a specific security policy.
127. **C.** Security must be included as an initial aspect of product design; it shouldn't be added after initial development.
128. **B.** A buffer overflow occurs if the operating system or the software fails to properly set boundaries and restrictions on how much data can be sent to the CPU.
129. **B.** Nonvolatile storage is labeled as secondary storage.
130. **C.** ROM is also known as firmware.
131. **D.** The CPU is the most trusted physical computer component because it is the central element of a system. All the other components are controlled by or accessed from the CPU.
132. **C.** Software is not trusted so virtual memory is used to create an access control layer between software and the physical components of the computer (that is, the kernel and its resource managers, such as the virtual memory manager).
133. **C.** Ring 2 is designated for I/O device drivers.
134. **B.** The lower the number, the greater the protection provided by that ring.
135. **A.** A problem state is the state in which an application or problem is executing; it has nothing to do with errors.
136. **C.** Multitasking is processing more than one process at once.
137. **A.** The statement "The more complex a security system is, the less assurance it provides" is true.
- Protection can occur at any point between the subject and object. Security measures often regulate activities between programs and objects. The simpler the security system is, the more likely it will provide the intended security.
138. **B.** TCB is the collection of components within a system that provides a specific level of trust (that is, security).
139. **B.** Resource isolation is required to provide accountability on a system.
140. **B.** The Take-Grant model is represented by a directed graph that specifies the rights a subject can transfer to an object or that a subject can obtain from another subject.
141. **C.** A state machine requires secure transactions.
142. **C.** The rows of an access matrix are known as capability lists.
143. **D.** The Bell-LaPadula model protects the confidentiality of data.
144. **A.** The Biba model is lattice based.
145. **B.** The Clark-Wilson model is primarily concerned with the prevention of unauthorized modification of data.
146. **B.** The Clark-Wilson model requires separation of duties.
147. **D.** The Orange Book contains the details on Trusted Computer System Evaluation Criteria (TCSEC).
148. **B.** Common Criteria (CC) is a replacement for and update to TCSEC.
149. **A.** A is the highest security valuation as defined by TCSEC.
150. **B.** A B3 TCSEC certification requires the use of security domains.

660 PRACTICE EXAM





151. **C.** B2 is the highest TCSEC security designation that still allows for the presence of covert channels.
152. **B.** The NIACAP Site Accreditation type is used to evaluate a specific self-contained location.
153. **A.** A closed system does not have published specifications.
154. **B.** Timing and storage are the two most common types of covert channels.
155. **A.** A network-based IDS would be ineffective against a host-based backdoor; therefore, a host-based IDS should be used.
156. **B.** Separation of duties specifies that no single person has complete access to or control over all the security mechanisms on a system.
157. **C.** Rotation of duties reduces collusion because multiple people will have the skills to review the activities within any specific job position and detect fraud or other crimes. It also forces the criminal element to involve more people in the conspiracy to keep things quiet because each time jobs are rotated, new individuals become capable of detecting the crime.
158. **A.** The primary goal of change management is to ensure that all changes made to a system do not result in reduced security.
159. **D.** Change management should provide for rollback to a previous version of the system.
160. **C.** Change evaluation in light of industry security standards is not an appropriate procedure in the process of change management.
161. **A.** Skills assessment exams are not part of personnel security management.
162. **B.** Owners must show due care to avoid full responsibility for a security breach.
163. **A.** The act of piggybacking is when a person walks through a secured doorway without self-authenticating immediately behind someone who performed proper self-authentication.
164. **C.** Data remnants are the elements of data remaining on media after it has been erased.
165. **C.** Security controls should be transparent to the user.
166. **C.** Monitoring and auditing don't directly prevent attacks. The results of monitoring and auditing can be used to select countermeasures to protect against future attacks.
167. **C.** Social engineering is the monitoring activity that obtains information simply by asking for it.
168. **B.** A clipping level is the level below which all normal activities occur; only events above this level should be suspect.
169. **A.** Clipping levels are ineffective against slow, low-profile intrusion attempts.
170. **D.** Audit logs can be stored on removable media, but it is not a universal requirement.
171. **D.** Program designers including omission errors in their custom scripts is a threat because of accidental loss, not inappropriate activities.
172. **A.** The use of encryption does not prevent traffic analysis.
173. **B.** SATAN is a vulnerability scanner.
174. **C.** A strong password policy, although a good security measure, is not a countermeasure against port mapping. Useful port mapping countermeasures include filtering traffic at the firewall, disabling banners on network services, and deploying an IDS.

175. **B.** A superzapping program is designed to recover from a system freeze or malfunction by bypassing security and access controls.
176. **C.** A sniffer's ability to decode is used to reveal the contents of captured traffic.
177. **D.** nmap is a port scanner.
178. **D.** IPSec authentication is a countermeasure for session hijacking.
179. **B.** Business continuity planning is used to label or describe a minor disruptive event where an organization must recover and continue to support critical functions.
180. **A.** Upgrading security mechanisms is not a factor of business continuity planning. All the other selections are aspects or factors of business continuity planning.
181. **C.** Personnel safety is always the highest priority.
182. **D.** Senior management is ultimately responsible for the success of a business continuity plan.
183. **C.** Vulnerability assessment is often part of performing a BIA, but it is not one of the goals of a BIA.
184. **B.** Employee awareness is a key element in the implementation process of a business continuity plan. Senior management approval is not a key element because it's the step before implementation.
185. **A.** There should be only a single business continuity plan per organization.
186. **A.** Paying dividends is not an issue to be included in a disaster recovery plan.
187. **B.** A warm site has a functional facility with hardware but no software or configuration.
188. **A.** Hot sites have a high administrative overhead.
189. **C.** A cold site is not considered an adequate resource for disaster recovery because of the time required to install and configure systems for productive operation.
190. **C.** The distance from the original site is the most important aspect to consider. It should be far enough away not to be involved in the same disaster as the primary site but close enough that traveling is not extensive.
191. **B.** Electronic vaulting is the process of backing up data to an offsite location.
192. **C.** Remote journaling is parallel processing of transactions to an alternative site.
193. **D.** If a plan is not tested, it does not work.
194. **A.** A checklist test involves the distribution of the plan to all appropriate personnel for review.
195. **C.** A parallel test is a full test, but the activities at the production environment are not stopped.
196. **C.** A simulation test is a type of test that works through the recovery plan up to the point just before alternative processing is initiated.
197. **B.** The salvage team returns to the original site only after the possibility of personal injury is eliminated.
198. **C.** Only when the organization has fully returned to the original site is the emergency over.
199. **A.** The first step in recovering from a disaster should be the restoration of the least critical functions. This allows for testing of procedures, connectivity, infrastructure, and so on so that if there are any errors or problems, they can be detected and resolved before the critical functions of the organization are affected.

662 PRACTICE EXAM

200. **D.** Having a mechanism to continue to pay employees even if business production is stopped is an important and often overlooked aspect of disaster recovery planning.
201. **B.** Writing viruses is not specifically restricted in the (ISC)² Code of Ethics.
202. **C.** Selling products over the Internet is not considered an unethical activity by the IAB according to RFC 1087.
203. **D.** The GASSP does state that system security is bound by societal restraints or factors.
204. **A.** Wasting resources is not considered a computer crime.
205. **C.** TEMPEST is used to prevent the interception of RF emanations.
206. **B.** Masquerading is the act of pretending to be someone else to gain a greater level of access.
207. **A.** A salami attack is the theft of small amounts of information from numerous sources to reveal or extract highly confidential information.
208. **C.** Data diddling is the act of modifying data through unauthorized means.
209. **B.** Evidence gathering requires special skills, usually those of a systems expert or forensic specialist.
210. **A.** The severity of punishment is related to the degree the organization demonstrates due diligence.
211. **C.** Revealing the results of periodic vulnerability assessments is not part of due care.
212. **D.** Legal liability exists if the countermeasure is less than the expected loss from an exploited vulnerability.
213. **D.** The prudent man rule from the 1991 U.S. Federal Sentencing Guidelines states that senior officials must perform their duties with the same care that ordinary sensible people would exercise under similar circumstances.
214. **B.** A legally recognized obligation must be demonstrated to prove negligence.
215. **D.** Common law is based on precedent (in other words, court and judicial decisions established in previous cases).
216. **B.** Criminal law is directed toward protecting the public.
217. **A.** A patent provides the creator of a work exclusive rights for 17 years.
218. **B.** European privacy laws are more restrictive than U.S. privacy laws. For example, collecting personal data to use as marketing demographics is more strictly regulated in Europe than in the U.S.
219. **D.** Privacy can't be guaranteed when electronic monitoring is used.
220. **A.** Entrapment is the act of encouraging the commission of a crime by an individual who initially had no intention of committing a crime.
221. **D.** The CIRT team should retain and protect evidence, not purge it.
222. **B.** Imaging the hard drive is the only action out of this list of options that should be taken during the initial process of evidence gathering at the scene of a computer crime.
223. **C.** Sufficiency is not an aspect of evidence; that is up to a judge or jury.
224. **B.** Printouts should be labeled using permanent markers.
225. **B.** Prevention of alteration or tampering of evidence is the most important aspect of evidence gathering.

226. **A.** Direct evidence proves or disproves a specific act through oral testimony based on evidence gathered through the witness's five senses.
227. **C.** Hearsay evidence is not based on personal, firsthand knowledge of the witness but is obtained from another source.
228. **A.** Hearsay evidence is generally inadmissible in court.
229. **C.** Interrogation is the act of gathering or discovering enough evidence about a subject to consider an individual a suspect.
230. **D.** The U.S. Privacy Act of 1974 applies to federal agencies and is directed toward the protection of information about private individuals that is stored in government databases.
231. **C.** Training is an administrative security control.
232. **D.** Fire detection and suppression are technical security controls.
233. **B.** Property tax rate is not a security concern.
234. **D.** Telephone company proximity is not a security or safety consideration.
235. **D.** A human incompatible server/equipment area does not provide for or double as a personnel shelter.
236. **B.** Unauthorized disclosure violates confidentiality, not availability.
237. **C.** Security controls must always comply with laws and regulations.
238. **B.** Hardware should be replaced as it reaches its mean time between failures.
239. **B.** Sag is momentary low voltage.
240. **A.** Traverse mode noise is the EMI generated by the difference between hot and neutral wires.
241. **B.** 40%–60% humidity is ideal for the operation of computer components.
242. **C.** Static electricity of 2,000 volts will cause a system shutdown.
243. **C.** A Class C fire extinguisher should be used for electrical fires. Class A fire extinguishers are used for common combustibles. Class B fire extinguishers are used for liquid fires. There is no Class AB fire extinguisher.
244. **D.** A preaction pipe is recommended for computer centers because it can be disabled and drained in the event of a false alarm or quickly averted emergency before damaging electronic components.
245. **B.** FM-200 is the replacement gas for Halon.
246. **B.** Guards can't be used in numerous environments, and many environments don't support human presence or intervention.
247. **B.** Dogs are reliable perimeter controls.
248. **A.** A mantrap is a double set of doors often monitored by a guard.
249. **D.** Lighting is the most common form of perimeter or boundary protection.
250. **B.** Purging is the act of removing data remnants from media for use outside the protected environment.






P A R T

III

APPENDIXES

- A Glossary**
- B Overview of the Certification Process**
- C What's on the CD-ROM**
- D Using the *PrepLogic Practice Tests, Preview Edition Software***



Glossary

A

abstraction When data is managed as a collection called an object, it is called abstraction.

access control An extension of administrative procedures that tell administrators how to configure authentication and other access control features of the various components.

Address Resolution Protocol (ARP) Allows a host to determine an unknown remote destination physical address from a known logical address. It is typically used for mapping IP addresses to MAC addresses.

administrative management The management of all things administrative, such as personnel management, recordkeeping, and the like.

administrative or management controls Personnel screening, separation of duties, rotation of duties, and least privilege are examples of administrative controls.

American Standard Code for Information Interchange (ASCII) ACSII is most commonly used for text file formatting. ASCII uses a 7-bit binary number to represent characters.

Annual Loss Expectancy (ALE) A mathematical formula used in risk analysis to determine the potential amount of money represented by a business interruption event.

annualized rate of occurrence The ratio of the estimated possibility that a threat will take place in a one-year time frame.

application software maintenance controls These controls monitor installations, updates to applications, and changes.

Application Specific Integrated Circuit (ASIC) ASICs are special purpose computer chips that are designed to perform specific tasks and functions—for example, switching functions.

ARCnet This network access methodology uses a token-bus access method for delivering data at 2.5Mbps.

asset valuation The evaluation of assets and the risk associated with their loss.

assurance The confidence that a product or process meets security objectives defined for it.

Asynchronous Transfer Mode (ATM) ATM is a LAN/WAN transmission method that uses fixed length 53-byte cells for transmitting data at rates up to 10Gbps. ATM uses permanent virtual circuits and switched virtual circuits to identify connections.

audit An examination of a set of data against a set of rules to determine whether it is in compliance with the rules.

audit and variance detection controls Audit logs contain information on the exercise of privilege or records of system activity. Variance detection products detect and may send alerts when unusual activities occur.

authentication Authentication is a matter of what the entity knows, what they may have, or who the entity is. For strong authentication, use at least two of these principles.

668 Appendix A GLOSSARY

authenticity The requirement in law that evidence must be established as being authentic before it is accepted in court.

authorization The process of granting permission to specific resources.

awareness training Making employees aware of the importance of information security, its significance, and the specific security-related requirements relative to their position; the importance of confidentiality, proprietary, and private information.

B

backup procedures The procedures that detail how copies of data are kept so that they will be available should recovery be necessary. They should also address the potential need for equipment.

banner grabbing A technique in which Telnet or other sessions are started with a computer in hopes of getting the banners, or blurbs, which tell about the service, back for analysis. Banners can tell an attacker much information about the system.

baselines Used to create a minimum level of security necessary to meet policy requirements.

basic input/output system (BIOS) Provides the basic information on hardware devices including storage devices, as well as security, boot sequence.

Bell-LaPadula model Security policy model of the Orange book. It is a state transition model of security policy, and it describes access control rules. In this model, entities in a computer system are divided into an abstract set of subjects and objects; each change in computer system state must not change security. System state is secure if only access by subjects to objects is in accordance with policy. Policy grants clearance (is access authorized by this subject?) to a subject based on classification of the object.

Best Evidence Rule A requirement in law that evidence of a writing must normally be the original writing itself rather than a copy. The rule has many exceptions and is of little relevance to electronic evidence.

Biba Model integrity model Another formal access control mode. In this mode a set of rules states that a subject can't depend on object or other subject that is less trusted than itself.

blended malware Malware that can use several attack vectors to infect systems and networks. It also uses various techniques to do harm.

boot sector virus A virus that infects the boot sector of a computer.

Bootstrap Protocol (BootP) BootP is a protocol that allows for the automatic network configuration and booting of devices, particularly diskless workstations. BootP is a predecessor of DHCP.

bridge A data link layer network device that is used to segment network traffic. Bridges can learn the MAC addresses of hosts on segments which allows it to filter traffic from segments that do not contain the destination.

British Naval Connector (BNC) BNC connectors are used to connect coaxial networks using a half locking mechanism.

broadcast A broadcast is a packet or frame that is addressed to all hosts on a network.

brute-force attack An attack in which every possible combination of characters is tried in order to crack a password.

buffer overflow An error condition where too much data is entered into a program or some portion of a program. A buffer, or area in memory, is reserved to hold the entry and is too small for the amount of data entered. The result of a buffer overflow can be a simple crash of the program, or it can result in a situation where an attacker can run code of his choice on the system.

business continuity planning (BCP) The process of determining those critical business functions that must be quickly restored after a business interruption event if the business is to survive. Also, the development of steps to ensure this occurs. It encompasses both disaster-recovery planning and business-resumption planning.

business impact assessment (BIA) An analysis of the impact of the loss of business processes. A financial loss is calculated over time and used to determine the maximum tolerable downtime for each process. The BIA results are then used to identify the most critical processes and how quickly they must be brought back online. Resources can then be allocated to assist planners and business process owners in ensuring this activity.

business resumption planning (BRP) The process of detailing the recovery of critical operational processes.

C

cache CPU memory storage that the CPU can access faster than RAM. Level-2 cache is usually a dedicated, small memory subsystem, while Level-1 cache is a smaller memory subsystem that is built into the CPU chip.

capture The file of captured packets collected by the sniffer.

carrier sense, multiple access/collision detection (CSMA/CD) CSMA/CD is the network access methodology employed by Ethernet. With CSMA/CD, when a host decides to transmit, it first listens to determine whether it detects a signal. If it does not, it then attempts to transmit. Finally, it listens to determine whether a collision occurred and the data needs to be retransmitted.

catastrophe An event which causes enough damage to require significant restructuring of an environment.

centralized controlled computing Computers may be distributed but configuration, maintenance and control is centralized.

centralized system All computing takes place in one place.

chain of evidence (or Chain of Custody) A series of records showing where evidence came from, who was responsible for it, what happened to it, how it was protected, whether it was changed, and so on.

change control Maintenance and tracking of changes to hardware and software.

channel The path used for information system transfer.

channel service unit/data service unit (CSU/DSU) The CSU/DSU acts as a buffer between the CPE and the provider network, ensuring that faulty CPE cannot affect the provider network. The CSU/DSU converts data from LAN technologies to WAN technologies.

Clark-Wilson model of security policy An access control model designed for commercial deployment. It features nondiscretionary access control, privilege separation, and least privilege.

clearance A level associated with a user in a system that has mandatory access control. A user with a clearance can access information with a sensitivity label equal to or lower than her clearance.

clearing If writable media is to be reused, it is made available by overwriting the classified information. (This does not lower the classification level of the media.)

clipping level That level at which repeated errors will trigger an alert.

closed system A computer system that does not use normal user interfaces and limits users to a single application or language.

670 Appendix A GLOSSARY

cold site An alternative process site that only provides the basic environment. Wiring, power, and air conditioning should be available, but no computers or peripherals are present.

co-location A second location for business operations. Data is constantly refreshed at the co-location so that if the prime site fails, the co-location site can immediately take over operations. Web sites are often co-located to ensure constant and consistent operation no matter the interruption.

compartmentalization Isolation of OS, user programs, and data files from each other provides protection against unauthorized access. Also, breakdown of sensitive data into small blocks to reduce risk of unauthorized access.

computer facility The facilities where computers will be used, including the structures or parts of structures. For small computers, standalone systems, and word processing equipment, it may be defined as the physical area where the computer is used.

computer incident response team (CIRT) The CIRT is the group of people designated to respond to security incidents. CIRT is synonymous with CERT (Computer Emergency Response Team), but CERT is a trademark.

computer premises equipment (CPE) CPE refers to the customer-owned, -managed, and -maintained equipment at the customer location that typically connects to a service provider.

confidentiality The secrecy of the information asset.

confidentiality, integrity, and availability (CIA) Represents the three basic principles of computer security.

configuration management Maintenance and tracking of changes to hardware and software.

conflict of interest An unethical state of affairs in which a professional has incentive to serve two inconsistent objectives, such as a duty to serve her employer while she is being paid a bribe to serve a vendor to her employer.

controls The means to prevent misuse or abuse of privileges while allowing authorized individual or processes to do their jobs.

cooperative hot site A site owned by a group (departments, divisions within a company, partner companies, strategically aligned companies, or associations) and available to members of the group during an emergency.

copyright The exclusive right to exploit a written work such as a novel, photograph, or software program.

corrective control A control that reduces the impact of an attack.

counteranalysis A technique that seeks to confuse the enemy with misinformation.

countermeasure A method that will prevent or mitigate the effect of an attack.

covert channel Communications channel that allows information to be transferred outside of the security policy through an abnormal path which is therefore not protected by normal security.

covert storage channel Allows one process to store and another to read, from the same location. Each process has separate and different security levels.

covert timing channel One process signals another by modifying systems resource use, in order to affect the response time. The second process can see this difference.

cryptographic keys Using public key cryptography, the user has a private key, or digital signature, that is used to sign a common hash value that is sent to the authentication server. The server can then use the known public key for the user to decrypt the hash.

cyclic redundancy check (CRC) CRC is a mathematical calculation for ensuring data integrity. When the source system transmits a data frame, it calculates the CRC and places the result at the end of the frame.

When the destination receives the data frame, it recalculates the CRC and compares the result to the result that the source sent. If they match, the data is complete and error free. If they do not, there is an error in the data and it is discarded.

D

data classification The classification used is dependent on the overall sensitivity of the data and the levels of confidentiality desired.

data communication equipment (DCE) DCE is any device that connects a system to a communications channel or public network.

data consistency Data viewed or retrieved in different ways will be the same. A transaction will maintain data consistency.

data duplexing The process of data mirroring where two disk controllers are present. Data mirroring might also exist when only one disk controller is available but might be less efficient because the controller must be responsible for two disk writes.

data hiding This is when data is unknown by and inaccessible from other layers.

data independence A characteristic of database systems. The data stored in the database can be used by multiple applications, even by applications which have not been developed yet.

data mining An analysis technique that requires specialized software and highly trained analysts.

data mirroring The process of writing data twice. A minimum of two data drives is provided and data is written to both drives. Should one drive fail, the other can be used instead.

data recovery In the event of an error, or system crash, the system can recover. Transactions in process at the time of the crash are checked and either rolled back or forwarded to complete a transaction and maintain data consistency.

data redundancy The same data stored in multiple places.

data remanence Data left over after data is deleted from the system.

data reuse Data gathered for one use is made available elsewhere.

data terminal equipment (DTE) DTE is the system that connects to a communications channel or public network.

data-vaulting The process of storing data at remote locations by electronically moving the data. As data is modified at the prime location, it is refreshed at another location.

data warehouse An aggregate of an organization's information.

database management system (DMBS) The management processes control database creation, manipulation, and access.

decentralized Computing facilities exist throughout the company. They may or may not be linked with each other.

degauss To use a demagnetizer to alter the magnetic composition of the data media. This effectively cleans the disk leaving little trace. In short, the data cannot be recovered and the disk is reusable. In technical terms, a variable, alternating current (AC) field (in which current alternates from zero to some maximum value and back again) is applied for the purpose of demagnetizing magnetic recording media.

672 Appendix A GLOSSARY

degausser An electrical device (AC or DC), or a magnet assembly that can be used to degauss magnetic media.

denial of service (DoS) An attack on a computer system that results in legitimate users not being able to access it.

dense wave division multiplexing (DWDM) DWDM uses different colors, and thus wavelengths, of light to transmit multiple data streams simultaneously over a single physical connection.

destruction Physically altering ADP-system media or components so they are no longer usable for data storage or retrieval.

detective control A control that protects vulnerability, reduces impact of attack, or prevents its success.

deterrent control A control that reduces the likelihood of attack.

dictionary attack An attack on passwords that use the password encryption algorithm to encrypt each word in a dictionary and compare it to passwords in the encrypted password file. A match means a password has been found.

differential backup Data files that have changed since the last backup are copied during differential backups. Files are not marked as backed up. The next backup copies files changed since the differential backup, as well as all files previously copied in the differential backup. This continues until a full backup or incremental backup is performed.

Digital Millennium Copyright Act (DMCA) A federal law that makes it a crime to make, sell, or distribute products or services intended to circumvent the encryption or other technical devices that copyright owners use to protect their copyrighted material. The DMCA also makes it a crime to break encryption or other devices for the purpose of gaining unauthorized access to copyrighted material.

Directive on Data Protection A law within the European Union requiring the protection of personal information and forbidding the exportation of personal data to countries with inadequate privacy laws.

disaster recovery planning (DRP) The process of detailing the recovery of critical technology operations.

discretionary access control (DAC) Restricts access to system objects (files, directories, devices) based on user id and groups. A user with some access permission can pass this on to another user.

discretionary security protection In this model, users process data at their security level. Security features prevent over writing of system memory, or of interfering with other users' work.

distributed In a distributed environment, computers are everywhere and so is the processing of data.

dynamic random access memory (DRAM) Memory composed of transistors and paired capacitors.

E

eavesdropping The gathering of information by observing and listening in on transmitted data, for example with a sniffer.

elevated privileges attack An attack in which an attacker hopes to obtain or increase his privileges on a victim computer.

encryption Encryption uses algorithms to convert data into an unintelligible form. In basic terms, encryption uses a secret key, a private value, to perform a mathematical function on the data to make it unusable by the casual observer.

environment The collection of all circumstances, conditions, and objects, including external ones that have an effect on system development, operation, and maintenance.

erasure Magnetic media is expunged by degaussing, either by AC current or DC current or by using a magnet.

escort An appropriated cleared individual assigned to control the activities of the person begin escorted. The escort should have appropriate clearance and authorization as well as understand the security implications of the access and activities of the escorted person.

Ethernet A network protocol and cabling scheme that uses the CSMA/CD access method to transmit data at speeds from 10Mbps to 10Gbps.

ethical hacking A technique that uses hacker tools and techniques to attack a network or computer with the purpose of finding vulnerabilities and making them known to the owners of the network or computer.

Evaluation Assurance Level (EAL) Assurance components representing a point on the predefined assurance scale.

Exclusionary Rule A rule in constitutional law that aims to enforce the rights granted under the Fourth Amendment. The rule states that if evidence is collected in violation of the Fourth Amendment, that evidence shall be excluded from evidence in a trial, such as the trial of a suspected criminal.

export of labeled information Writing information to another system, while still maintaining the protection mechanism associated with it. This can be done by either by assigning security levels to output devices or by writing sensitive label with data.

exposure factor The frequency of event occurrence is used to estimate the percentage of loss on a particular asset because of a threat.

Extended Binary Coded Decimal Interchange Code (EBCDIC) EBCDIC is a proprietary IBM method for encoding characters in an 8-bit binary number.

extranet An extranet is a network connection that provides external access to internal resources. Extranets typically refer to the connection between communications partners networks.

F

fail-over cluster Multiple processors, drives, and other hardware work together to provide an environment where the failure of one component (CPU, drive, and so on) will not mean the failure of processing. Should one system fail, the other takes up the operation.

fair information practices Recognized methods for protecting privacy of personal data. They include the rights of the data subject to notice about how data will be collected and used, choice about whether it will be collected, and reasonable protection of the data to ensure accuracy, integrity, and security.

Federal Emergency Management Agency (FEMA) A U.S. agency charged with providing support and funding during and after disasters.

fiber distributed data interface (FDDI) FDDI is a token-passing ring methodology that uses dual rings to deliver data at 100Mbps.

fiber optic Fiber optic describes a cable type that uses discrete pulses of light over specially manufactured optical cables for the transmission of data. Fiber optic cable is not susceptible to electro-magnetic interference.

File Transfer Protocol (FTP) FTP provides for the transfer of files using a client/server model.

firewall A firewall is a perimeter security device that is designed to filter unwanted traffic from reaching protected resources. Firewalls act as points of entry to protected networks.

flooding net A collection of compromised machines that are used by an attacker to attack some other victim.

674 Appendix A GLOSSARY

forensics The use of science and technology to investigate and establish facts that can be used in court.

formal security model A mathematically precise statement of security policy. The model gives the initial state of system and notes the process by which the system progresses from one state to another. It defines what is meant by a definition of secure state of system. This statement should be supported by formal proof: If the initial state of system satisfies the definition of *secure*, all future states will be secure.

formal verification With formal verification, an automated tool is used to design and test a highly trusted system. It demonstrates the following features: design consistency between a formal specification and a formal security policy model, and implementation consistency between formal specifications and high-level program implementation.

Fourth Amendment Part of the U.S. Constitution that guarantees citizens protection from unreasonable searches and seizures by the government.

frame relay Frame relay is a WAN switching technique that uses virtual circuits and bandwidth on demand for the transmission of data.

full backup A complete copy of all data on the disk is performed.

full duplex Full duplex is the ability to transmit and receive at the same time.

full recovery test The process of testing all aspects of recovery.

G

gateway A gateway is an entry point to or from a network. Gateways are often routers or firewalls. Gateways can be used to provide access between networks using different technologies and protocols.

Gauss A unit measure representing the magnetic flux density produced by a magnet or other magnetizing force.

goods Materials and supplies including inspection and test equipment. Technical data is not included.

Gramm-Leach-Bliley Also known as the Financial Services Modernization Act, which requires financial institutions to give consumers notice about how personal information about them will be used. It also requires institutions to implement safeguards to protect personally identifiable information.

grid computing The combination of the excess capacity of all computers on a network to perform additional processing.

guidelines Recommendations for how policies can be implemented.

H

half-duplex Half-duplex is the capability to transmit or receive, but to only be able to perform one operation at a time.

hardware segmentation The isolation of software processes and data via the separation of hardware.

Healthcare Insurance Portability and Accountability Act (HIPAA) The Act generally requires healthcare providers to maintain the confidentiality of patient information.

hearsay An out-of-court statement that is being offered as evidence in court. Evidence law often prohibits hearsay from being used in court.

hierarchical database Data is organized in a tree structure with a tree being composed of branches or nodes. Think of the branches as if they are data records, the leaves of the branches are the data. (One example of a hierarchical database is IMS.)

hierarchical storage management (HSM) The dynamic and automatic management of the storage and retrieval of online data files.

high-level data link control (HDLC) HDLC is a data link-layer bit-oriented synchronous protocol that is typically used for providing WAN connectivity.

high-speed serial interface (HSSI) HSSI is a point-to-point protocol that defines transmission speeds of up to 52MBps over short distances. HSSI is often used to connect to ATM and T3 connections.

host-based intrusion detection system (HIDS) A program that runs on servers and workstations to detect intrusions against the host.

hot site An alternative site that is completely configured with equipment, systems software, and an appropriate operating environment. It is only necessary to provide personnel, programs, and data.

hub A hub is a layer-1 device that functions as a multiport repeater. Hubs do not look at or verify the data, but rather they simply receive, boost, and retransmit signals.

hybrid site Some combination of hot, cold, or warm sites.

incident response Procedures that discuss how to involve management in the response as well as when to involve law enforcement.

incremental backup Copies data files that have changed since the last backup. Backed-up files are marked and the next backup will not include these files.

indicator Information that may be seen, heard, or collected from Web sites, tapes, discs, documents, and observations.

information label A label that is associated with a subject or object (such as a file). It is similar to sensitivity labels, but different, because sensitivity labels may have classification, categories, and dissemination markings, and handling caveats (EYES ONLY). Information labels can change as information content of subject or object changes, while sensitivity labels remain static.

initial program load (IPL) The start-up process of a mainframe.

Institute of Electrical and Electronics Engineers (IEEE) The IEEE acts as a coordinating and governing body handling networking, computing, and communications standards.

integrity The assurance that the data is accurate and reliable.

Integrated Services Digital Network (ISDN) A technology that was designed to transmit digital data over existing telephone networks.

International Standards Organization (ISO) An international standards making body that is responsible for defining global standards for communications and data exchange.

Internet The connection of networks that provides connectivity between networks and resources on a global basis.

Internet Control Message Protocol (ICMP) ICMP is used on IP networks to provide error reporting, management, and control information.

Internet facing A computer or device that has a direct connection to the Internet.

Internet Package Exchange (IPX) IPX is a Novell-proprietary network layer protocol that is used for transmitting data across a network.

Internet Protocol (IP) IP is a layer-3 protocol that defines the logical addressing of hosts using IP addresses. IP also provides for the routing of data by the use of network identifiers as a part of IP addresses.

676 Appendix A GLOSSARY

Internet Relay Chat (IRC) IRC is a real-time client/server protocol that allows hosts to communicate with each other interactively.

intranet An intranet is modeled on the Internet and refers to the private design of a network that transmits data for internal use only.

intrusion detection A methodology for determining if a system is under attack.

intrusion detection systems (IDS) Software or hardware devices that are programmed to analyze network traffic or system logs and to raise an alert if it detects potentially hazardous network traffic and programs that would indicate an attack or intrusion is taking place.

intrusion prevention A methodology to prevent the new and unknown attacks and enforcement of application behavior.

intrusion prevention systems (IPS) Software or hardware that is designed to prevent desktops or servers from being exploited by new and unknown attacks.

IPSec A TCP/IP security protocol. It offers authentication of network devices, port filtering, integrity, and encryption.

J–K

knowledge-based systems Often called expert systems, these attempt to parallel the thought process and deduction effort that transpires when an expert searches for the answer to a problem.

L

labeling In a mandatory access control system, it is the requirement to assign sensitivity labels to every subject or object in a system.

Layer-2 Tunneling Protocol (L2TP) An encryption/tunneling protocol designed to provide temporary secure channels of communications across the internet.

layer-3 switch A network device that has routing and switching capabilities built into the device, reducing the need for multiple devices to perform the tasks.

layering Where the system resources are managed in a protected kernel and everything else runs in an outer layer known as *user's space*. If a process running in the user's space wants to access a protected resource, such as the disk, it makes a request to the kernel layer in order to perform the action.

Link Access Protocol-Balanced (LAPB) A WAN specification that defines the communications between DCE and DTE.

license A contract for the right to use property, such as copyrighted software.

linear bus A linear bus defines the connection of multiple devices to single cable in a linear fashion.

linear printer daemon (LPD) LPD provides for remote printing capabilities to network attached print devices using TCP/IP.

local area network (LAN) A collection of network devices that are able to share resources and communicate with each other.

lock-and-key protection system An access control or protection system that requires matching a key or password with a specific access requirement.

logic bomb Code that is designed to execute because of some event, such as a calculation result or day of the year.

M

macro virus A virus written using the macro language present in desktop applications such as Word or Excel. The macro language enables users of these applications to automate repetitive tasks, such as opening multiple files. The macro automatically executes when the file is opened. Virus writers take advantage of this to write a malicious macro—for example, one that deletes a file. One way the virus writer can then cause harm is to send the now infected, but seemingly harmless, document to the user. The user opens the document and suffers the harm.

magnetic field intensity (MFI) MFI represent the magnetic force required to produce a desired magnetic flux. It is symbolized in an equation with the letter H .

magnetic flux Lines of force representing a magnetic field.

magnetic flux density (MFD) MFD represents the strength of a magnetic field. It is symbolized as the letter B (see also Gauss).

magnetic remanence After a magnetic force is applied, some magnetic flux density will remain. This represents data that remains on the media after degaussing.

magnetic saturation Magnetic saturation is the amount of magnetizing force, in which the most magnetic flux will occur. Increasing the magnetizing force produces little increase in magnetic flux.

malware Programs written to do harm.

mandatory access control (MAC) Restricts access to objects based on sensitivity of information in object, object label, and authorization of subject (clearance). Mandatory system enforces users can't share their files.

maximum tolerable downtime (MTD) The amount of time that a business process can be non-operational and the business can still survive.

media Physical components, such as tape reels, floppy diskettes, hard drives, and so on, used for data storage.

mesh A network topology in which all devices are connected to every other device, providing complete redundancy.

mirror image An identical copy of the data on a hard disk. It is better to conduct forensic analysis on the copy than on the original data on the hard disk.

mobile site A facility that exists in trailers and therefore can be moved to a location near the existing facility.

modem A contraction of Modulator/Demodulator. A modem provides for the signaling conversion between analog and digital systems.

multicast A multicast is an addressing method in which data is delivered to multiple hosts, but not to all hosts.

multilevel device Non-removable drive, capable of inheriting sensitivity labels, so user can't just copy data to an untrusted system or device.

multi-partite virus Infects files, boot sector, and master boot records.

multiplexer A multiplexer is a device that merges multiple low-speed transmissions into a single high-speed channel. A device at the remote end reverses the process, breaking the single transmission back into the individual transmissions.

N

Netware Core Protocol (NCP) NCP is the Presentation layer protocol that translated data into a format that can be understood by Novell Netware networks.

network database (IDMS/R) data is represented in blocks or record types. Blocks include data fields. Arrows between the blocks represent a relationship between the data.

678 Appendix A GLOSSARY

network address translation (NAT) NAT is the translation of addresses on one network to addresses on another. It is typically used to translate from internal to public addresses.

network file system (NFS) NFS is a UDP-based file sharing mechanism, typically used for Unix-based networks.

network interface card (NIC) A NIC is a piece of hardware that provides network access to a host system.

network intrusion detection system (NIDS) A NIDS is used to detect unauthorized or malicious data on network segments.

Network News Transfer Protocol (NNTP) NNTP is a network protocol for defining the posting, retrieval and management of data to newsgroups.

non-essential records Records that are not critical for business continuity. They can be easily recovered or replaced.

nonrepudiation The ability to ensure the authenticity of a message by verifying it is using the message's digital signature.

N-type Connector N-type connectors are screw together connectors that are typically used for interconnecting thicknet/10base5 cabling.

O

object data model A model in which data in an application is associated with a central entity. For example, an object "person" includes all the associated data that defines the person, including address, telephone number, position, and supervisor. In the object model, methods or functions the object can do are also associated with the object. In our "person" object, methods might be "change password" or "change address."

object-oriented database Combines the object data model of object-oriented programming with DBMS.

object-oriented programming A programming model in which an object data model is used.

Oersted A unit of measure which represents the necessary magnetizing force which will produce the desired magnetic flux across a surface.

open storage The condition where classified information is stored in an accredited facility, but is not GSA-approved secure containers, nor are authorized personnel in the facility.

open system A computer system that uses normal user interfaces and provides total system access to the user.

Open Systems Interconnect (OSI) OSI is a reference model that is used to define the processes that must occur to enable network communications.

operational controls Operational controls protect day-to-day procedures and include mechanisms such as physical and environmental protection, privileged entry commands, backup, contingency planning, documentation, change control management, hardware controls, and input and output controls.

OPSEC Process The process of understanding your day-to-day operations from the viewpoint of a competitor, enemy, or hacker and then developing and applying countermeasures.

Orange book The common name for the first United States official government security specification. The book was so named because of its orange color.

overwrite See overwrite procedure.

overwrite procedure A procedure which makes unreadable data or destroys data on a writable storage media by recording patterns of unclassified data over or on top of the data stored on the media.

overwriting See overwrite procedure.

P

packet A short block of data that is transmitted on a network.

packet analyzer A program or device that is able to capture and analyze different types of data traffic on a network.

parallel test A test of recovery procedures where the objective is to perform processing equivalent to a complete business cycle. The test is run separately from normal operations.

partial backup Only changed data is copied. Incremental and differential backups are examples of partial backups.

password The most common form of authentication.

patent The exclusive right to exploit a unique invention.

pen test Shorthand for penetration testing.

penetration testing The testing of network security. This is done by using common hacker tools and methodologies in an effort to find vulnerabilities. Countermeasures, such as patches, configuration, or workarounds can then be used to harden security.

Perimeter Intrusion Detection and Assessment System (PIDAS) Often in the form of a fence equipped with various sensors.

persistence Data remains the same after code is executed.

physical control space (PCS) The spherical space surrounding information processing (electronic) equipment. This space, expressed in meters, should be under enough physical control to prevent hostile intercept of emanations. PCS can be controlled by fences, guards, patrols, walls, and so forth depending on resources available.

physical safeguards Items such as fire suppressant systems, alarms, and power backup or conditioning which are made available in order to mitigate a disaster.

physical security Physical protection which is provided for resources against deliberate and accidental physical threats.

Point-to-Point Protocol (PPP) PPP is used to transmit data over serial or dial-up point-to-point connections, giving the appearance that the remote host is just another node on the network.

Point-to-Point Tunneling Protocol (PPTP) PPTP is a Microsoft developed protocol that provides for VPN connections between hosts and networks.

policies Standards and guidelines that will be used throughout your organization to maintain your security posture.

polymorphic virus A virus that changes its own code to evade detection.

port redirection tool A tool that allows an attacker to use an open port on a firewall to access a target and then attack an entirely different port.

port scanner A program that attempts to determine whether any of a range of ports is open on a particular computer or device.

Post Office Protocol 3 (POP3) POP3 provides for incoming message storage and retrieval of email messages.

primary key The column in a database table which is selected for the primary index. It allows a relationship to be built with other tables that include the same column.

primary storage Another word for main memory. It is volatile or temporary memory, otherwise known as Random Access Memory (RAM). (Disks are referred to as secondary storage.)

680 Appendix A GLOSSARY

privacy The protection of personally identifiable information from corruption or unauthorized access.

Privacy Enhanced Mail (PEM) PEM is a proprietary RSA encryption method for ensuring the privacy of email messages.

privilege The right to do something on a computer such as log on, add users to a group, backup files, and so on.

privileged instruction Instructions that only the operating system can run. This code may also address areas of memory or other components restricted to the OS. The OS must be running in supervisor or kernel mode to use these instructions.

procedural safeguards Processes such as safety inspections, fire drills, and security awareness training that will mitigate the effects of a disaster, or perhaps prevent it from occurring.

procedures Mechanisms put into place to ensure the integrity of information and to prevent attacks on the storage of that data (*contamination*) and on its transmission (*interference*).

process isolation The ability to run different processes on one computer and yet separate them from one another. Each process has its own data and code space. Consequently, if a process fails, it can only crash itself; other running processes are not affected.

promiscuous mode An operational mode of a network interface card that changes the normal behavior of the card from only listen to information addressed to it, to one where the card listens for all traffic on the network.

protection profiles Implementation-independent set of security requirements for the category of Target Of Evaluation (TOEs) that meet a selection need.

protocol analyzers A type of sniffer.

proxy A proxy is a device that filters requests between systems. Proxies intercept data and make the requests on behalf of the source system.

purging The orderly removal of obsolete data files and data by erasure, overwriting of storage or resetting of registers.

Q

qualitative risk analysis Estimated loss is used to evaluate the risk.

quantitative risk analysis A mathematical approach to risk analysis in which the probability of occurrence is multiplied times the calculated monetary loss.

R

random access Also known as direct access. Some index, or other capability, exists that allows a search to go directly to the record required.

rapid application development A software development method that uses focus groups, prototyping, and a shortened timeframe.

real memory The Random Access Memory provided by the system hardware.

recovery point objective (RPO) The goal for restoring a business process.

recovery time objective (RTO) The amount of time available to restore a critical business process.

redundant array of inexpensive disks (RAID) RAID provides for fault tolerance of data by using redundant disks for the storage of either mirrored data or parity data that can be used to re-create the original data.

redundant site An alternative site that exactly mirrors the current data processing environment.

reference monitor An abstract machine that enforces TOE access control policies.

referential data integrity The database rule that says no database record can refer to the primary key of a non-existent table.

registers High-speed memory locations in the CPU. There are only a few of these locations.

relational database Data is stored in tables that consist of rows (like records in a regular file) and columns (like fields). Relationships are formed between tables based on a selected primary key.

remanence Remanence may be used to indicate the data left on storage media after the power is turned off. It is also a measure of the magnetic flux density that remains on media after degaussing.

remote authentication dial-in user service (RADIUS) RADIUS is a protocol that provides for the authentication of remote connections and users to network resources.

remote procedure call (RPC) RPC is a client/server architecture that is used for distributed programming.

repeater A repeater is a network device that simply boosts and retransmits signals without reading any of the data being transmitted. Repeaters function at the physical layer.

restricted area An area secured by restrictions and controls in order to safeguard property or material.

Reverse Address Resolution Protocol (RARP) RARP is very similar to ARP; however RARP resolves known MAC addresses to unknown IP addresses.

revision control The maintenance and tracking of changes to hardware and software.

ring A ring is a network topology in which devices are interconnected to each other in a circular fashion.

ring zero The inner core of the operating system. When the computer is running, different code is said to run at different levels. Ring zero is reserved for privileged instructions and access by the operating system itself.

risk analysis The process of determining if a threat is likely to occur and if it does, what damage will occur.

risk management The identification, measurement, control, and minimization of loss associated with uncertain events or risks.

router A router is a device which can deliver data to remote networks by using logical addresses and routing protocols to determine the path to the remote network.

S

Safe Harbor on Data Protection An arrangement between the European Union and the U.S. government under which U.S. companies can establish that they are complying with European privacy law by agreeing to protect personal data collected in Europe.

sanitization The elimination of classified information from magnetic media to permit the reuse of the media at a lower classification level or to permit the release to uncleared personnel or personnel without the proper information access authorizations.

Sanitized media Magnetic media that can be declassified after classified data is erased or overwritten.

secondary storage Nonvolatile storage. A variety of actual media that can store data and code for a very long time; includes devices such as disks, tapes, and CD-ROMs.

secure electronic transmission (SET) SET was developed to provide a framework for protecting the use of credit cards used in Internet transactions against fraud by using PKI to ensure data integrity and authentication.

682 Appendix A GLOSSARY**Secure/Multipurpose Internet Mail Extension**

(S/MIME) S/MIME is an email security standard that uses RSA public key exchanges.

Secure Sockets Layer (SSL) SSL provides for Transport layer encryption authentication and data integrity.

secure state None of the subjects can access objects in an unauthorized manner.

security area A physically defined space that contains classified matter (documents or material), which are subject to physical protection and personnel access controls.

security control architecture The sum of controls built into a system.

security controls A database provides variable security controls by limiting access to those who require it.

security function Part of the Target Of Evaluation (TOE) that enforces a subset of rules.

security kernel Hardware, firmware, and software that implement the reference monitor. The security channel must be complete (it mediates all access), be isolated (protected from modification) and be verifiable (can be verified as correct).

security level Sensitivity of information, from a sensitivity label.

security model Precise statement of security rule of a system.

security perimeter The boundary of security controls.

security target A set of requirements and specifications that are used as a base for evaluating a Target of Evaluation (TOE).

segmentation Hardware protection features, virtual memory is divided in segments, process may use many segments, unprivileged user processes cannot access or modify memory used by system.

sensitive information That which, if disclosed, altered, lost, or destroyed, could adversely affect national security or other federal government interest.

sequential access Data is searched by starting at the beginning of the media or file and searching every bit of data until the requested information is found.

Serial Line Internet Protocol (SLIP) SLIP is used to provide temporary network connections over telephone networks using IP.

server message block (SMB) SMB is the presentation layer protocol responsible for translating data into a format that Microsoft networking recognizes.

shunt trip A switch that can be used to immediately shut off power to a location.

Simple Key Management for Internet Protocol (SKIP) SKIP is a stateless network layer encryption mechanism developed and used primarily for SUN Solaris environments, though it functions on Windows-based systems as well. SKIP is able to encrypt data without needing a prior message exchange between hosts in order to establish a secure channel. Consequently, SKIP can be used in simplex communications environments.

Simple Message Transfer Protocol (SMTP) SMTP is a protocol that is used to deliver email messages to remote servers.

single-level device Support for sensitivity level, which is dependent on physical location or inherent level of security of device type. (For example, workstations, printers, communication ports, removable media.)

single-loss expectancy The amount of the potential loss for a specific threat.

sniffers Devices or software programs that capture packets and decode them.

spoofing An attack technique where some characteristic is misrepresented. An IP source spoof means the IP address of another system is inserted in a packet to replace the source address of the attacker's system.

star property (also known as **property or confinement property*) Bell-LaPadula security model rule that allows a subject write access to an object if the security level of the subject is dominated by the security level of object.

static random access memory Level 2 cache, usually consists of several transistors but no capacitor.

storage area network (SAN) Storage area networks that are centrally managed and network accessible storage systems.

Structured Query Language (SQL) SQL is a protocol that defines the formatting of data for use in mainframes and database communications.

structured walkthrough test A test in which members of the team walk through the plan looking for and correcting weaknesses.

supervisor or kernel mode The opposite of User mode. Supervisor mode is the mode within which the OS runs.

survivability The capability of a system to continue to process critical applications in spite of the fact that it suffered disruptive or damaging events (such as contamination with dust, an earthquake, a bomb, and so on).

swIPe swIPe is a predecessor to IPSec. swIPe provides encryption at the network layer by encapsulating the original packet within the swIPe packet. swIPe does not have policy or key management functionality built into the protocol.

switch A switch is a data link device that can filter, forward or flood traffic based on MAC address, thereby reducing contention in a network.

switched multimegabit data service (SMDS)

SMDS is a high-speed packet switching technology for use over public networks. It is provided for companies that need to send and receive large amounts of data on a bursty basis, providing for connectionless communications. It is a bandwidth-on-demand technology.

switched networks Networks in which switches are used to deliver packets from one computer to another. The switch forms a connection between the devices on the fly and no computer is exposed to traffic from every computer on the network.

synchronous Synchronous refers to the clocking or timing of data transmissions.

synchronous data link control (SDLC) SDLC is a bit-oriented, synchronous protocol that is typically used for interconnectivity between IBM SNA devices.

system development life cycle The series of steps that tracks the development of applications, from concept through disposal.

system downtime The time when the system is purposefully shut down or made unavailable in order to perform maintenance.

system outage The system is unavailable due to some non-planned event.

T

target of evaluation (TOE) IT product, system, and associate administrator and user guidance documentation—the subject of an evaluation.

technical controls Audit and journaling, integrity validations such as checksums, authentication and file system permissions.

Telnet Telnet is an application-layer protocol that provides for remote terminal emulation capabilities for TCP/IP-based hosts.

684 Appendix A GLOSSARY

Terminal Access Controller Access Control System Plus (TACACS+) TACACS+ is a remote authentication protocol. Although it has a similar function to RADIUS, TACACS+ differentiates itself by separating the authentication and authorization capabilities, as well as using TCP for connectivity. As a result, TACACS+ is generally regarded as being more reliable than RADIUS.

threat A person, event, or thing which has the ability to cause harm along with the intention to do so.

Time of Check to Time of Use (TOC/TOU) If an instruction is executed in more than one step, it may be possible to compromise the system by attacking between the steps.

tip-off indicator An indicator that provides focus for the attacker. They tell him where to concentrate his efforts.

TOE Security Functions (TSF) The combination of hardware, software, and firmware of the TOE. It enforces the TOE Security Policy.

TOE Security Policy (TSP) The set of rules which determine how TOE assets are managed, and protected.

token A form of one-time password authentication that satisfies the “what you have” scenario.

token ring Token ring refers to a network access methodology that uses a token-passing access method over a ring topology to transmit at speeds of 4MBps–16MBps.

trade secret The right to exclusive use of confidential commercial information.

Transmission Control Protocol (TCP) TCP is a transport-layer protocol that provides for reliable data delivery and connection-oriented communications.

Transmission Control Protocol/Internet Protocol (TCP/IP) TCP/IP is a suite of protocols that defines network communications governing media access, packet transport, session communications, and application functions.

Transport Layer Security (TLS) TLS is a Transport layer security mechanism that provides for encryption of data and access authentication.

trap door Portals that circumvent system protection. They are often legitimate debugging techniques that are accidentally or purposefully left in production code.

Trivial File Transfer Protocol (TFTP) TFTP is a subset of FTP that provides for the transfer of files without authentication. TFTP is a UDP-based transmission method.

Trojan horse A program that masquerades as something else in order to trick a user into running it.

trusted channel A means whereby a remote IT product and TSF can communicate.

Trusted Computer Security Evaluation Criteria (TCSEC) Also referred to as the Orange book of the rainbow series, TCSEC was developed by the Department of Defense to provide guidelines for evaluating vendor security.

trusted computing base (TCB) The sum of hardware, software, and firmware that enforces a security policy for a product. A TCB can enforce a security policy if it contains the appropriate mechanism and is correctly configured by the administrator.

trusted distribution The movement of trusted systems from vendor to customer, in exact evaluated system shipped by vendor.

trusted facility management Assures separation of duties, operator, administer, security administrator, with duties clearly defined for each role.

Trusted Network Interpretation (TNI) The TNI is referred to as the “Redbook” of the rainbow series. The TNI, or Redbook, interprets the TCSEC.

trusted path User communicated directly with Trusted Computing base. Can't be initiated by untrusted software. With a trusted path, no software can mimic trusted software.

trusted system A system developed in accordance with orange book criteria and evaluated by these criteria.

tunnel A tunnel is the encapsulation of one protocol within another, often providing security and encryption of the original data.

type-1 magnetic media Magnetic media with coercivity factors not exceeding 325 Oersteds.

type-2 magnetic media Magnetic media with coercivity factors exceeding 325 Oersteds, possibly as high as 750 Oersteds (also known as high-energy media).

U

unicast A unicast is an addressing method in which data is addressed to a specific host.

unshielded twisted pair (UTP) UTP is a point-to-point cable type that provides for both voice and data grade transmissions.

User Datagram Protocol (UDP) UDP is a Transport layer protocol that provides for unreliable, connectionless communications.

user mode The mode in which applications and other instructions used by ordinary operators, or individuals are run.

V

validation Tests and evaluates to determine if security specs and requirements are met.

verification Compares two levels of specification to ensure correspondence between them.

verify backup The process whereby a backup system checks a tape backup to ensure it is viable.

virtual local area network (VLAN) A VLAN is the logical separation of systems over a physically connected network. VLANs are generally synonymous with subnets.

virtual memory The combination of real memory and that provided by disk paging or swap files.

virtual private network (VPN) A VPN provides for secure transmission of data over an otherwise insecure medium by encrypting the data in a tunnel.

virus A program loaded onto a computer without the permission of the owner and then run without permission. Viral code hides itself within legitimate code.

vital records Records that have critical importance to the company and whose loss or damage would have critical impact on business continuity.

vulnerability A weakness in a computer system, software, device, infrastructure or operation which may allow a threat to succeed.

W

war dialer A program or device that automatically dials a range of phone numbers and reports on those that are answered by a computer or fax machine.

warm site This alternative site might be partially configured. Some peripheral equipment, such as printers might be available.

Web services Small, reusable programs that can be accessed from otherwise unconnected sources. Web services may be written in XML and used to communicate across the Internet or an organization's intranet.

worm Malware that spreads itself from one computer to another across a network.

X-Z

X Window A remote graphical user interface emulation protocol that is typically used for Unix connectivity. Similar in concept to Telnet, X Window provides for the remote display of the GUI environment.

X.25 X.25 is a highly reliable WAN connection technique that functions at the physical and data link layers of the OSI model. X.25 uses virtual circuits for establishing the communications channel between hosts.

X.400 X.400 is a messaging formatting standard that defines how addressing is performed.

xDSL An acronym for multiple types of Digital Subscriber Line. xDSL is a high bandwidth broadband connection method that is typically used for Small Office and Home Office connectivity.

Overview of the Certification Process

This appendix explains the CISSP certification process and looks at what is involved in taking the CISSP Exam. At the time of writing, this information is accurate; however, (ISC)² reserves the right to change exam and certification track information at any time, thus, it's worth checking their Web site at <http://www.ISC2.org> to see whether there have been any changes to the program.

DESCRIPTION OF THE PATH TO CERTIFICATION

The CISSP examination process is distinct from the CISSP certification process. You must pass the exam to obtain the certification, but passing the exam is no guarantee that you will achieve certification.

To sit the exam you must do the following:

- ◆ Pay the fee
- ◆ Assert that you have the years of experience required (until Jan. 1, 2003, it's three years, and after that date it's four years or three years of experience plus a college degree)
- ◆ Complete a candidate agreement that includes a legal obligation to adhere to the code of ethics and asserts the truth of the experience statement given

- ◆ Answer questions regarding criminal history and background

To obtain the certification you must do the following:

- ◆ Pass the exam with a score of 700 or more
- ◆ Complete an endorsement form that provides validation by a CISSP or an officer of your corporation, which attests to your experience

ABOUT THE CERTIFICATION PROGRAM

You might be asking why this exam is for you, and why now? Besides the fact that the certification brings certain obvious professional benefits to you, the CISSP program gives you access to the (ISC)² organization and the benefits that access affords. In addition, the CISSP exam is well recognized in the Infosec community:

- ◆ **Recognized proof of professional achievement**—This is a level of competence that is commonly accepted and valued by the industry.

688 Appendix B OVERVIEW OF THE CERTIFICATION PROCESS

- ◆ **Enhanced job opportunities**—Many employers give preference in hiring to applicants who have certification. They view certification as proof that a new hire knows the procedures and technologies required.
- ◆ **Opportunity for advancement**—Certification can be a plus when an employer awards job advancements and promotions.
- ◆ **Training requirement**—Certification might be required as a prerequisite to attending a vendor's training course, so employers often offer advanced training to employees who are already certified.
- ◆ **Customer confidence**—As the general public learns about certification, customers will require that only certified technicians be assigned to their accounts.

For any additional information or clarification about the CISSP certification path and its history and benefits, consult the (ISC)² home page at www.isc2.org. As discussed earlier, you can also check this site to see whether there have been any recent changes in the certification program.

What's on the CD-ROM

This appendix is a brief rundown of what you'll find on the CD-ROM that comes with this book. For a more detailed description of the *PrepLogic Practice Tests, Preview Edition* exam simulation software, see Appendix D, "Using the *PrepLogic Practice Tests, Preview Edition* Software." In addition to the *PrepLogic Practice Tests, Preview Edition*, the CD-ROM includes the electronic version of the book in Portable Document Format (PDF), several utility and application programs, and a complete listing of test objectives and where they are covered in the book.

PREPLOGIC PRACTICE TESTS, PREVIEW EDITION

PrepLogic is a leading provider of certification training tools. Trusted by certification students worldwide, we believe PrepLogic is the best practice exam software available. In addition to providing a means of evaluating your knowledge of the Training Guide material, *PrepLogic Practice Tests, Preview Edition* features several innovations that help you to improve your mastery of the subject matter.




For example, the practice tests allow you to check your score by exam area or domain to determine which topics you need to study more. Another feature allows you to obtain immediate feedback on your responses in the form of explanations for the correct and incorrect answers.

PrepLogic Practice Tests, Preview Edition exhibits most of the full functionality of the *Premium Edition* but offers only a fraction of the total questions. To get the complete set of practice questions and exam functionality, visit PrepLogic.com and order the *Premium Edition* for this and other challenging exam titles.

Again for a more detailed description of the *PrepLogic Practice Tests, Preview Edition* features, see Appendix D.

EXCLUSIVE ELECTRONIC VERSION OF TEXT

The CD-ROM also contains the electronic version of this book in PDF. This electronic version comes complete with all figures as they appear in the book. You will find that the search capabilities of the reader comes in handy for study and review purposes.



APPENDIX **D**

Using the *PrepLogic Practice Tests, Preview Edition* Software

This Training Guide includes a special version of PrepLogic Practice Tests—a revolutionary test engine designed to give you the best in certification exam preparation. PrepLogic offers sample and practice exams for many of today's most in-demand and challenging technical certifications. This special Preview Edition is included with this book as a tool to use in assessing your knowledge of the Training Guide material while also providing you with the experience of taking an electronic exam.

This appendix describes in detail what *PrepLogic Practice Tests, Preview Edition* is, how it works, and what it can do to help you prepare for the exam. Note that although the Preview Edition includes all the test simulation functions of the complete, retail version, it contains only a single practice test. The Premium Edition, available at PrepLogic.com, contains the complete set of challenging practice exams designed to optimize your learning experience.

EXAM SIMULATION

One of the main functions of *PrepLogic Practice Tests, Preview Edition* is exam simulation. To prepare you to take the actual vendor certification exam, PrepLogic is designed to offer the most effective exam simulation available.

Question Quality

The questions provided in the *PrepLogic Practice Tests, Preview Edition* are written to highest standards of technical accuracy. The questions tap the content of the Training Guide chapters and help you review and assess your knowledge before you take the actual exam.

Interface Design

The *PrepLogic Practice Tests, Preview Edition* exam simulation interface provides you with the experience of taking an electronic exam. This enables you to effectively prepare you for taking the actual exam by making the test experience a familiar one. Using this test simulation can help eliminate the sense of surprise or anxiety you might experience in the testing center because you will already be acquainted with computerized testing.

Effective Learning Environment

The *PrepLogic Practice Tests, Preview Edition* interface provides a learning environment that not only tests you through the computer, but also teaches the material you need to know to pass the certification exam.

692 Appendix D USING THE PREPLOGIC PRACTICE TESTS, PREVIEW EDITION SOFTWARE

Each question comes with a detailed explanation of the correct answer and often provides reasons the other options are incorrect. This information helps to reinforce the knowledge you already have and also provides practical information you can use on the job.

- ◆ The Installation Wizard copies the *PrepLogic Practice Tests, Preview Edition* files to your hard drive; adds *PrepLogic Practice Tests, Preview Edition* to your Desktop and Program menu; and installs test engine components to the appropriate system folders.

SOFTWARE REQUIREMENTS

PrepLogic Practice Tests requires a computer with the following:

- ◆ Microsoft Windows 98, Windows Me, Windows NT 4.0, Windows 2000, or Windows XP.
- ◆ A 166MHz or faster processor is recommended.
- ◆ A minimum of 32MB of RAM.
- ◆ As with any Windows application, the more memory, the better your performance.
- ◆ 10MB of hard drive space.

Installing PrepLogic Practice Tests, Preview Edition

Install *PrepLogic Practice Tests, Preview Edition* by running the setup program on the *PrepLogic Practice Tests, Preview Edition* CD. Follow these instructions to install the software on your computer:

- ◆ Insert the CD into your CD-ROM drive. The Autorun feature of Windows should launch the software. If you have Autorun disabled, click Start and select Run. Go to the root directory of the CD and select setup.exe. Click Open, and then click OK.

Removing PrepLogic Practice Tests, Preview Edition from Your Computer

If you elect to remove the *PrepLogic Practice Tests, Preview Edition* product from your computer, an uninstall process has been included to ensure that it is removed from your system safely and completely. Follow these instructions to remove *PrepLogic Practice Tests, Preview Edition* from your computer:

- ◆ Select Start, Settings, Control Panel.
- ◆ Double-click the Add/Remove Programs icon.
- ◆ You are presented with a list of software installed on your computer. Select the appropriate *PrepLogic Practice Tests, Preview Edition* title you want to remove. Click the Add/Remove button. The software is then removed from your computer.

USING PREPLOGIC PRACTICE TESTS, PREVIEW EDITION

PrepLogic is designed to be user friendly and intuitive. Because the software has a smooth learning curve, your time is maximized because you start practicing almost immediately. *PrepLogic Practice Tests, Preview Edition* has two major modes of study: Practice Test and Flash Review.

Using Practice Test mode, you can develop your test-taking abilities as well as your knowledge through the use of the Show Answer option. While you are taking the test, you can expose the answers along with a detailed explanation of why the given answers are right or wrong. This gives you the ability to better understand the material presented.

Flash Review is designed to reinforce exam topics rather than quiz you. In this mode, you will be shown a series of questions but no answer choices. Instead, you will be given a button that reveals the correct answer to the question and a full explanation for that answer.

Starting a Practice Test Mode Session

Practice Test mode enables you to control the exam experience in ways that actual certification exams do not allow:

- ◆ **Enable Show Answer Button**—Activates the Show Answer button allowing you to view the correct answer(s) and full explanation for each question during the exam. When not enabled, you must wait until after your exam has been graded to view the correct answer(s) and explanation.
- ◆ **Enable Item Review Button**—Activates the Item Review button allowing you to view your answer choices, marked questions, and facilitating navigation between questions.

To begin studying in Practice Test mode, click the Practice Test radio button from the main exam customization screen. This will enable the options detailed previously.

To your left, you are presented with the option of selecting the preconfigured Practice Test or creating your own Custom Test. The preconfigured test has a fixed time limit and number of questions. Custom Tests allow you to configure the time limit and the number of questions in your exam.

The Preview Edition included with this book includes a single preconfigured Practice Test. Get the complete set of challenging PrepLogic Practice Tests at PrepLogic.com and make certain you're ready for the big exam.

Click the Begin Exam button to begin your exam.

Starting a Flash Review Mode Session

Flash Review mode provides you with an easy way to reinforce topics covered in the practice questions. To begin studying in Flash Review mode, click the Flash Review radio button from the main exam customization screen. Select either the preconfigured Practice Test or create your own Custom Test.

Click the Best Exam button to begin your Flash Review of the exam questions.

Standard PrepLogic Practice Tests, Preview Edition Options

The following list describes the function of each of the buttons you see. Depending on the options, some of the buttons will be grayed out and inaccessible or missing completely. Buttons that are appropriate are active. The buttons are as follows:

- ◆ **Exhibit**—This button is visible if an exhibit is provided to support the question. An exhibit is an image that provides supplemental information necessary to answer the question.

- ◆ **Item Review**—This button leaves the question window and opens the Item Review screen. From this screen you will see all questions, your answers, and your marked items. You will also see correct answers listed here when appropriate.
- ◆ **Show Answer**—This option displays the correct answer with an explanation of why it is correct. If you select this option, the current question is not scored.
- ◆ **Mark Item**—Check this box to tag a question you need to review further. You can view and navigate your Marked Items by clicking the Item Review button (if enabled). When grading your exam, you will be notified if you have marked items remaining.
- ◆ **Previous Item**—View the previous question.
- ◆ **Next Item**—View the next question.
- ◆ **Grade Exam**—When you have completed your exam, click to end your exam and view your detailed score report. If you have unanswered or marked items remaining, you will be asked if you would like to continue taking your exam or view your exam report.

Time Remaining

If the test is timed, the time remaining is displayed on the upper-right corner of the application screen. It counts down minutes and seconds remaining to complete the test. If you run out of time, you will be asked if you want to continue taking the test or if you want to end your exam.

Your Examination Score Report

The Examination Score Report screen appears when the Practice Test mode ends—as the result of time expiration, completion of all questions, or your decision to terminate early.

This screen provides you with a graphical display of your test score with a breakdown of scores by topic domain. The graphical display at the top of the screen compares your overall score with the PrepLogic Exam Competency Score.

The PrepLogic Exam Competency Score reflects the level of subject competency required to pass this vendor's exam. While this score does not directly translate to a passing score, consistently matching or exceeding this score does suggest you possess the knowledge to pass the actual vendor exam.

Review Your Exam

From Your Score Report screen, you can review the exam that you just completed by clicking on the View Items button. Navigate through the items viewing the questions, your answers, the correct answers, and the explanations for those questions. You can return to your score report by clicking the View Items button.

Get More Exams

Each *PrepLogic Practice Tests, Preview Edition* that accompanies your training guide contains a single PrepLogic Practice Test. Certification students worldwide trust PrepLogic Practice Tests to help them pass their IT certification exams the first time. Purchase the Premium Edition of PrepLogic Practice Tests and get the entire set of all new challenging Practice Tests for this exam. PrepLogic Practice Tests—Because You Want to Pass the First Time.

CONTACTING PREPLOGIC

If you would like to contact PrepLogic for any reason including information about our extensive line of certification practice tests, we invite you to do so. Please contact us online at www.preplogic.com.

Customer Service

If you have a damaged product and need a replacement or refund, please call the following phone number:

800-858-7674




Product Suggestions and Comments

We value your input! Please email your suggestions and comments to the following address:

feedback@preplogic.com

LICENSE AGREEMENT

YOU MUST AGREE TO THE TERMS AND CONDITIONS OUTLINED IN THE END USER LICENSE AGREEMENT (“EULA”) PRESENTED TO YOU DURING THE INSTALLATION PROCESS. IF YOU DO NOT AGREE TO THESE TERMS, DO NOT INSTALL THE SOFTWARE.



Index

SYMBOLS

* (star) property

Bell-LaPadula security model, 31-32, 345

Biba security model, 32

3DES. *See* Triple-DES

5-4-3 rule, 579

8mm tape, 157

10BASE-2 networks, 79, 579

10BASE-5 networks, 79, 579

200Mbps Fast Ethernet, 97

A

A division (Orange Book), 134

A1 class (Orange Book), 134, 358

abstraction, 217, 351, 587, 598

acceptable usage policy (AUP), 214, 224

access control. *See also* ACLs (access control lists)

administration, 27-29, 576

attacks

brute-force attacks, 41, 577

denial-of-service attacks, 42, 577

dictionary attacks, 41

sniffing attacks, 43, 578

spoofing attacks, 42-43, 578

authentication. *See* authentication

Bell-LaPadula model, 30-33, 576

Biba model, 32-33, 577

case study, 52-53

centralized access control, 38, 577

decentralized access control, 38-40, 577

defining data access, 209

discretionary access control, 20, 576

exam objective overview, 13-14, 17

identification. *See* identification

IPSec standard, 371

lattice-based access control, 22-25, 576

Liptner's lattice, 33, 577

mandatory access control, 21-22, 576

noninference models, 33, 577

penetration testing. *See* penetration testing

physical access controls, 540-544, 618

procedures, 211

reference monitor, 348-349

remote access. *See* remote access

role-based access control, 26-27, 576

rule-based access control, 25, 576

storage area networks (SANs), 260

versus accountability, 18-19

versus authentication, 17-18

access logs, 541-542

access servers, 119, 583

accidents, 447

accountability, 18-19, 188, 576, 586

logging, 19

accounts receivable, insurance coverage, 461

accreditation, 284

ACKs (acknowledgements), 127

ACLs (access control lists), 27, 347, 576, 597

versus labels, 353

ACM (Configuration Management) class, 367

698 INDEX

- active attacks, 45
- active monitoring, 98, 189
- active physical access controls, 541
- ActiveX applets, 246
- acts of war, 446
- actual cash value (ACV), 461
- Address Resolution Protocol (ARP), 130
- administrative controls, 389, 604
- administrative law, 497, 613
- administrative management (operations security), 418-420, 608
- administrative procedures, 211
- ADO (Delivery and Operation) class, 367
- ADSL (Asymmetric Digital Subscriber Line) connections, 117
- ADV (Development) class, 368
- AES (Advanced Encryption Standard), 314
- AGD (Guidance Documents) class, 368
- AH (Authentication Header), 373
- air conditioning, 547-548
- ALC (Life Cycle Support) class, 368
- ALE (annualized loss expectancy), 200-204, 408
- all risks, 461, 611
- allow mode (IPSec), 372
- AMA (Maintenance of Assurance) class, 368
- American Registry for Internet Numbers (ARIN), 407
- annualized loss expectancy (ALE), 200-204, 408
- annualized rate of occurrence (ARO), 200
- anomaly detection, 47
- anti-symmetric property (lattice-based access control), 23-25
- antidisaster procedures, creating, 468-469, 612
- antistatic mats, 546
- antistatic sprays, 546
- antiviral controls, 423-425, 609
- antivirus software, 277
- APE (Protection Profile Evaluation) class, 367
- application errors (risk factor), 193
- Application layer, OSI model, 72, 578
 - protocols, 126-127, 136-137
- application logs, 398
- application software maintenance controls, 389, 604
- application-filtering firewalls, 105, 581
- applications development. *See* systems development
- architecture, *See* security architecture
- ARCnet (Attached Resource Computer Network), 99
- ARIN (American Registry for Internet Numbers), 407
- ARO (annualized rate of occurrence), 200
- ARP (Address Resolution Protocol), 130
- Arrangement on the Recognition of Common Criteria Certificates in the Field of IT Security. *See* Common Criteria
- array indexing errors, 294
- ASE (Security Target Evaluation) class, 367
- assets
 - classification of, 533-535
 - vulnerabilities, 535-537
- asymmetric encryption, 217, 315-316, 595
- asynchronous tokens, 187
- Asynchronous Transfer Mode (ATM) networks, 114
- ATE (Tests) class, 368
- AtGuard, 402
- ATM (Asynchronous Transfer Mode) networks, 114
- Attached Resource Computer Network (ARCnet), 99
- Attachment Unit Interface (AUI) connections, 81
- attack signature recognition, 403
- attacks
 - active attacks, 45
 - birthday attacks, 325, 596
 - brute-force attacks, 41, 148, 266-267, 322-323, 577
 - chosen-ciphertext attacks, 322
 - chosen-plaintext attacks, 322
 - cipher-text only attacks (COAs), 321
 - collateral attacks, 412
 - corrective controls, 390
 - distributed denial-of-service (DDoS) attacks, 151, 269, 468

- denial-of-service (DoS) attacks, 42, 267-269, 577
 - detective controls, 390
 - deterrent controls, 390
 - dictionary attacks, 41, 266-267
 - elevated privileges attacks, 407
 - integrity attacks, 182
 - jump-point attacks, 412
 - keyboard attacks, 426
 - known versus unknown attacks, 45
 - known-plaintext attacks (KPA), 321-322
 - LAND attacks, 151
 - man-in-the-middle attacks, 323-324, 596
 - meet-in-the-middle attacks, 324, 596
 - NAK (negative acknowledgement) attacks, 272
 - network abuses
 - class A network abuses, 147-148
 - class B network abuses, 148-149
 - class C network abuses, 149-150
 - class D network abuses, 150-151
 - class E network abuses, 152-154
 - class F network abuses, 154-155
 - nuke attacks, 412
 - passive attacks, 45
 - preventative controls, 390
 - probing attacks, 154-155
 - pseudoflaw attacks, 272
 - random attacks, 412
 - replay attacks, 323, 596
 - as risk factor, 193
 - SMURF attacks, 151, 268
 - spoofing attacks, 42-43, 152, 269-270, 578
 - strategic attacks, 412
 - teardrop attacks, 151
 - Trojan horses, 152-153, 243, 247
 - Audit (FAU) class, 366**
 - audit controls, 389, 604**
 - audit logs, 542**
 - audit trails, 512**
 - auditing, 188-190, 395-398, 586, 605**
 - keystroke monitoring, 189-190
 - procedures, 211
 - protecting audit data, 190
 - AUI (Attachment Unit Interface) connections, 81**
 - AUP (acceptable usage policy), 224**
 - authentic evidence, 511**
 - authentication, 184-187, 311-312, 389, 577, 585-586**
 - biometrics, 36
 - case study, 52-53
 - MACs (message authentication codes), 316
 - passwords, 35-36, 185-187
 - RADIUS (Remote Authentication Dial-In User Service), 38
 - remote access authentication, 124
 - SSO (single sign-on) scheme, 37
 - SANs (storage area networks), 260
 - strong authentication, 185
 - TACACS (Terminal Access Controller Access Control System), 38
 - ticket schemes, 36-37
 - versus access control, 17-18
 - Authentication Header (AH), 373**
 - automatic sprinkler systems, 553**
 - AVA (Vulnerability Assessment) class, 368**
 - availability, 131, 183, 585**
 - awareness training, 227-228**
 - AXENT, 403**
- ## B
- B division (Orange Book), 133-134, 358**
 - B1 class (Orange Book), 133, 358**
 - B2 class (Orange Book), 134, 358**
 - B3 class (Orange Book), 134, 358**
 - back doors, 153**
 - Back Orifice, 273**

700 INDEX

- backups, 156-158, 472-481, 585**
 - alternative sites, 478-481, 612-613
 - alternatives to tape, 481
 - differential backups, 157, 475-476
 - full backups, 156, 474
 - hardware backups, 474, 478, 612
 - incremental backups, 157, 475-476
 - media, 158
 - partial backups, 474
 - tape storage, 477
 - vital records, 477
- backward chaining, 262**
- banner abuse, 154**
- banner grabbing, 407**
- banners, unauthorized access, 504-505**
- baselines (information security policies), 210**
- basic rate interface (BRI), 117**
- BCP. *See* business continuity planning**
- behavior-based intrusion detection systems, 141**
- Bell-LaPadula security model, 30-33, 216, 343-345, 576**
 - * (star) property, 31-32, 345
 - simple security rule, 30-31
 - subjects/principles versus users, 30
- best evidence rule, 511-512, 616**
- best practices**
 - operations security, 420-428
 - antiviral controls, 423-425
 - change management control, 427-428
 - privileged operation functions, 421-422
 - protecting sensitive information, 425-427
 - system development controls, 285
- BIA. *See* business impact assessment**
- Biba security model, 32-33, 345, 577**
- biometrics, 36, 543**
- BIOS (basic input/output system), 257**
- birthday attacks, 325, 596**
- BlackICE, 402**
- blended malware, 247**
- block mode (IPSec), 372**
- BLP. *See* Bell-LaPadula security model**
- blue beep (war dialer), 406**
- Bluetooth, 88**
- BMP (Bitmap) files, 73**
- BNC (British Naval Connector) connections, 81**
- Boehm, Barry, 281-282**
- boiler and machinery, insurance coverage, 461**
- BootP (Bootstrap Protocol), 126**
- bounds checking, 294**
- BRI (basic rate interface), 117**
- bridges, 76, 100-101, 581**
- broadcast storms, 100**
- broadcasts, 76-77, 94, 580**
- brownouts, 545, 618**
- brute-force attacks, 41, 148, 266-267, 322-323, 577**
- buffer (fiber-optic cable), 85**
- buffer overflows, 150-151, 293-294, 389**
- building construction, 540**
- Building Secure Software* (Viega and McGraw), 238, 293**
- Bureau of Export Administration (BXA), 502**
- bus topology, 89-91, 580**
- business continuity planning (BCP), 610**
 - BIA (business impact assessment), 452-454
 - e-commerce and, 455
 - gathering and charting information, 454-455
 - reporting, 458
 - validating the maximum tolerable downtime (MTD), 456-457
 - case study, 482-483
 - duplicated records, 460
 - emergency control centers, 459
 - FEMA checklists, 451
 - implementing the plan, 464
 - listing potential disasters, 445-448
 - maintaining the plan, 465-466
 - monetary losses, calculating, 454-455
 - objective overview, 439-440

operational plans, 458-459
 getting help, 460
 planning for insurance claim processing, 461-463
 providing item recovery details, 463-464
 reviewing insurance coverage, 460-461, 466
 reasons for having, 444-445
 risk analysis, 447-448
 scope of, determining, 451-452
 testing the plan, 464-465
 versus disaster recovery planning, 448-450

business impact assessment (BIA), 452-454, 610
 e-commerce and, 455
 gathering and charting information, 454-455
 reporting, 458
 validating the maximum tolerable downtime (MTD),
 456-457

business interruption insurance, 461

business-related computer attacks, 505

BXA (Bureau of Export Administration), 502

C

C division (Orange Book), 133, 358

C1 class (Orange Book), 133, 358

C2 class (Orange Book), 133, 358

CA (certificate authority), 188

cabling, 579-580
 coaxial, 79-82, 579
 failures, 158
 fiber-optic, 84-87, 580
 UTP (unshielded twisted pair), 82-84, 579-580

cache, 257

cameras, 561

**Canadian Trusted Computer Product Evaluation
 Criteria (CTCPEC)**, 356

capacitance detectors, 560

Cap'n Crunch, 275

captures, 399

carbon dioxide extinguishers, 551

cascading errors, 195

CASE (computer-aided software engineering), 291-292,
 594

CBA (cost-benefit analysis), 194-195, 203-204, 587

CBAC (Content-Based Access Control), 138

CC. *See* **Common Criteria**

CCDs (charge-coupled devices), 561

CD/DVD drives, 157

CDs (compact discs), 553-555

cell-switched connections, WANs, 114, 582

Center for Academic and Research Computing, 281-282

centralized access control, 38, 577

centrally controlled computing, 589

certificate authority (CA), 188

certification, 361

certification process, 687-688

CESG (Communications-Electronics Security Group),
 360

CFCs (chlorofluorocarbon compounds), 552

chain of evidence, 512-513, 616

challenge response schemes, 36

change control, 226-227, 389, 427-428, 609

charge-coupled devices (CCDs), 561

check points, 249-250

checksums, 389

chlorofluorocarbon compounds (CFCs), 552

chosen-ciphertext attacks, 322, 596

chosen-plaintext attacks, 322, 596

CIO Cyberthreat Response and Reporting Guidelines,
 507

ciphertext, 313, 595
 chosen-ciphertext attacks, 322, 596
 ciphertext-only attacks, 321

ciphertext-only attacks (COAs), 321

CIR (Committed Information Rate), 113

circuit proxy firewalls, 105

702 INDEX

- circuit-switched connections, WANs, 113, 582
- circumstantial evidence, 616
- CIRT (Computer Incident Response Team), 141-142
- Cisco Secure IDS, 402
- CISSP certification exam, 687-688
 - described, 623
- civil laws, 497, 613
- cladding (fiber-optic cable), 84
- Clark-Wilson security model, 346, 597
- class A network abuse, 147-148
- class B network abuse, 148-149
- class C network abuse, 149-150
- class D network abuse, 150-151
- class E network abuse, 152-154
- class F network abuse, 154-155
- classifying data, 218-222
 - commercial classification, 219
 - criteria for, 221
 - government classification, 220
 - procedures, 221-222
- CLEFs (Commercial Evaluation Facilities), 361
- client-based dial-in remote access, 119-120
- client-based VPNs (virtual private networks), 121-122
- clipping level, 189, 403
- closed systems, 350-351
- clustering, 262-263, 585
 - data clustering, 156
 - network services clustering, 156
- CMOS batteries, 258
- co-location, 468, 474, 612
- COAs (ciphertext-only attacks), 321
- coaxial cabling, 79-82, 579
 - failures, 158
- code libraries, 264
- Code of Ethics, (ISC)², 518-519
- CodeRed, 247, 406
- coding practices, 286, 292
 - computer-aided software engineering (CASE), 291-292, 594
 - malicious code, 274-277
 - object-oriented programming, 289-291, 594
 - structured programming, 286-289, 594
- cognitive passwords, 186
- cold sites, 478, 613
- collapsed backbones, 92
- collateral attacks, 412
- collisions, 76-77
- combination locks, 542
- commercial classification of data, 219
- Commercial Evaluation Facilities (CLEFs), 361
- Common Criteria (CC), 362-369, 601-603
 - areas not addressed by, 369
 - EALs (Evaluation Assurance Levels), 368-369, 602
 - objectives, 363
 - Part 1 (Introduction and General Model), 364-365
 - Part 2 (Security Functional Requirements), 365-367
 - Part 3 (Security Assurance Requirements), 367-368
 - versus other security standards, 370
- Communication (FCO) class, 366
- Communications-Electronics Security Group (CESG), 360
- compartmentation, 21-22
- compartmented security mode, 352, 598
- Computer Associates eTrust Intrusion Detection, 402
- computer centers, 540
- computer crime, 503-505, 615-616
 - security incidents
 - advance planning, 506-507
 - investigation of, 507-509
 - legal evidence, 510-513
 - major categories of, 505-506
- computer crime laws. *See* laws
- computer ethics, 517-519, 617

- computer forensics, 513-517
 - case study, 519-520
- Computer Fraud and Abuse Act, 503-504, 615
- Computer Incident Response Team (CIRT), 141-142
- Computer Security Handbook*, 455
- computer-aided software engineering (CASE), 291-292
- confidential data, 219-220, 588
- confidentiality, 181-182, 310-311, 585
 - access control lists, 347
 - Bell-LaPadula security model, 30-32, 343-345
 - IPSec standard, 371
 - networks and, 130-131
 - Orange Book standard. *See* Orange Book standard
 - SANs (storage area networks), 260
- configuration management, 226-227
- Configuration Management (ACM) class, 367
- configuration procedures, 211
- conflicts of interest, 519
- connections, network
 - cell-switched, 114, 582
 - circuit-switched, 113, 582
 - dedicated, 111-112
 - Frame Relay, 116, 159, 583
 - HSSI (High Speed Serial Interface), 118
 - ISDN (Integrated Services Data Network), 116-117, 583
 - packet-switched, 113, 582
 - SDLC (Synchronous Data-Link Control), 116, 583
 - X.25, 115, 583
 - xDSL (Digital Subscriber Line), 117-118, 583
- Constitution (U.S.), Fourth Amendment, 513, 616
- consumer fraud-related computer attacks, 505
- contamination, 182
- Content-Based Access Control (CBAC), 138
- contention, 90
- contention-based media access, 95
- controlled access protection (Orange Book, class C2), 358
- controlled security mode, 353, 599
- controls, 604
 - identifying available controls, 389-391
 - system development controls, 277-285
 - best practices, 285
 - RAD (Rapid Application Development), 282-283, 593
 - security control architecture, 283-285, 593
 - spiral lifecycle model, 280-282, 592-593
 - waterfall lifecycle model, 278-280, 592
- copyrights, 498-499, 613
 - case study, 520-521
- core (fiber-optic cable), 84
- core dumps, 558
- corrective controls, 390, 604
- cost-benefit analysis (CBA), 194-195, 203-204, 587
- counteranalysis, 395
- countermeasures, 393, 408
 - cost/benefit analysis, 203-204, 587
 - disgruntled employees, 415-416
 - employee-related threats, 412-414
 - hiring and firing/exit practices, 414-415
 - information system threats, 409-410
 - Internet-based threats, 416-417
 - mainframe threats, 410-411
 - physical threats, 417-418
 - threat risk analysis, 408-409
- covert storage channels, 354, 599
- covert timing channels, 354
- Crack program, 323
- crackers, 275
- credit card memory, 258
- criminal law, 497, 503-505, 613-615
- cryptanalysts, 320
- Cryptographic Support (FCS) class, 366

704 INDEX

cryptography, 181-182, 187, 594-595. *See also*

encryption

- asymmetric encryption, 315-316, 595
- authentication and, 311-312, 594
- confidentiality and, 310-311, 594
- digital signatures, 317, 595
- hash functions, 316-317, 595
- integrity and, 311, 594
- nonrepudiation and, 312
- objective overview, 307
- one-time ciphers, 318
- symmetric encryption, 313-314, 595

CSMA/CD (Carrier Sense, Multiple Access/Collision Detection), 95-96

CSU/DSU (Channel Service Unit/Data Service Unit), 119, 583

CTCPEC (Canadian Trusted Computer Product Evaluation Criteria), 356

Cybersafe, 403

CycSecure, 261

D

D division (Orange Book), 133, 358

DAC. *See* discretionary access control

daisy-chaining, 89

damage (risk category), 193

DAT (digital audio tape), 157

data access, defining, 209

Data Circuit-Terminating Equipment, 76

data classification, 218-222, 587-588

- commercial classification, 219

- criteria for, 221

- government classification, 220, 588

- procedures, 221-222

data clustering, 156, 585

data disposal, 556-559

Data Encryption Standard (DES), 218, 314

data hiding, 217, 587

data integrity. *See* integrity

data marts, 255

data mining, 255

data models, 251-252

data remanance, 427

data safes, 469

data storage, 256-259

- document libraries, 555-556

- electronic media, 553-555

- offsite, 559

- RAID (redundant array of inexpensive disks), 155-156, 474, 584-585

- SANs (storage area networks), 259-260

data striping with parity, 474

Data Terminal Equipment (DTE), 76

data vaulting, 468, 473-474, 612

data warehouses, 255

database management system (DBMS), 249-255, 589-590

databases, 249-255, 590-591

- data models, 251-252

- distributed databases, 252

- hierarchical, 252

- network databases, 252

- object-oriented databases, 252

- packed, 557

- relational databases, 251

- versus data marts, 255

- versus data warehouses, 255

Data Link layer, OSI model, 75-77, 579

DBMS (database management systems), 249-255, 589-590

DCE (Data Circuit-Terminating Equipment), 76

DDoS (distributed denial-of-service) attacks, 151, 269, 468

- decentralized access control, 38-40, 577
 - domains, 39-40
 - trust relationships, 40
- decryption, 313, 595
- dedicated connections, WANs, 111-112
- dedicated security mode, 352, 598
- degaussers, 557-558
- degaussing, 425-427, 557-558
- delayed loss, 196
- deleting user accounts, 28
- Delivery and Operation (ADO) class, 367
- demonstrative evidence, 510
- denial-of-service (DoS) attacks, 42, 267-269, 577
- dense wave division multiplexing (DWDM), 87
- DES (Data Encryption Standard), 218, 314
- destruction vulnerabilities, 535-536
- detective controls, 390, 604
- deterrent controls, 390, 604
- Development (ADV) class, 368
- dial-up access, 119-120
- dictionary attacks, 41, 266-267
- differential backups, 157, 475-476, 585
- digital audio tape (DAT), 157
- digital certificates, 319
- Digital Immune System for Cyberspace (IBM), 248
- digital linear tape (DLT), 157
- Digital Millennium Copyright Act (DMCA), 499
- digital signatures, 317, 595
 - hash functions, 316-317
 - nonrepudiation, 188, 312
- digital V-Ohm meters, 82
- direct access, 257
- direct evidence, 616
- directed broadcasts, 94
- Directive on Data Protection (European Union), 501
- dirty power supplies, 544
- disabling user accounts, 28
- disaster recovery planning (DRP), 466-467, 611-612.
 - See also* business continuity planning
 - antidisaster procedures, creating, 468-469
 - backups, 472-481
 - alternative sites, 478-481
 - alternatives to tape, 481
 - hardware backups, 478
 - procedures and policies, 474-477
 - tape storage, 477
 - vital records, 477
 - contact numbers, recording, 472
 - emergency control centers, 459
 - emergency response procedures, 470-471
 - listing potential disasters, 445-448
 - necessary resources, listing, 469-470
 - normal operations, restoring, 472
 - objective overview, 439-440
 - risk analysis, 447-448
 - scope of plan, determining, 468
 - step-by-step instructions, creating, 471
 - versus business continuity planning, 448-450
- disclosure (risk category), 193
 - vulnerabilities, 535-536
- discretionary access control, 20, 576
- discretionary security property (Bell-LaPadula security model), 345
- discretionary security protection (Orange Book, class C1), 358
- disks, 553-555
 - degaussing, 558
- disposing of data. *See* data disposal
- distilled water, 553
- distributed databases, 252, 590
- distributed denial-of-service (DDoS) attacks, 151, 269, 468
- distributed systems, 589
 - examples of, 244-245
 - malware for, 246-248
 - massively distributed systems, 245

706 INDEX

distribution cable, 85
 DIVX (Digital Video Express) files, 73
 DLT (digital linear tape), 157
 DMCA (Digital Millennium Copyright Act), 499
 document libraries, 555-556
 documentary evidence, 510
 documentation, 190-191
 domains, 39-40
 domain separation, 352, 598
 dominance relations (Bell-LaPadula security model), 344
 doors, 542-543
 DoS (denial-of-service) attacks, 42, 267-269, 577
 DRAM (dynamic RAM), 257
 Draper, John, 275
 DRP. *See* disaster recovery planning
 dry standpipe systems, 553
 DS-0 (Digital Signal Level 0) connections, 112
 DS-1 (Digital Signal Level 1) connections, 112
 DS-3 (Digital Signal Level 3) connections, 112
 DTE (Data Terminal Equipment), 76
 dual-homed host firewalls, 109
 due diligence, 211-212
 dumb cards, 542
 dumpster diving, 150, 556
 duplexing, 474
 DVDs, 553-555
 DWDM (dense wave division multiplexing), 87
 dynamic data storage, 480
 dynamic packet filtering firewalls, 106
 dynamic RAM (DRAM), 257

E

e-commerce, business impact assessment and, 455
 E1 connections, 112
 E3 connections, 112
 EAC (Expected Annual Cost), 408

EALs (Evaluation Assurance Levels), 368-369, 602
 edge servers, 277
 EF (exposure factor), 200
 Electronic Communications Privacy Act, 504
 elevated privileges attacks, 407
 email, 146-147
 abuse of, 149
 antiviral controls, 423-425
 gateways, 72
 protocols, 136-137
 emergency control centers, 459
 emergency power shut-off switches, 471
 employment policies, 222-225
 acceptable usage policy (AUP), 224
 background checks, 222-223
 employee-related threats, 412-414, 606-607
 employee theft, 537
 job descriptions, 225
 job rotation, 225
 termination, 224-225
 encapsulation, 77-78
 Encapsulation Security Payload (ESP), 373
 encryption, 181-182, 217-218, 587. *See also*
 cryptology
 asymmetric encryption, 315-316, 595
 birthday attacks, 325, 596
 brute-force attacks, 322-323, 596
 case study, 325-327
 chosen-ciphertext attacks, 322, 596
 chosen-plaintext attacks, 322, 596
 ciphertext-only attacks, 321, 596
 key length, 317-318
 known-plaintext attacks, 321-322, 596
 man-in-the-middle attacks, 323-324, 596
 meet-in-the-middle attacks, 324, 596
 replay attacks, 323, 596
 strong encryption, 321
 symmetric encryption, 313-314, 595

environmental and physical procedures, 211
 environmental controls, 426, 547-548, 619
 environmental spills, 447
 erasing data, 556-559
 ESP (Encapsulation Security Payload), 373
 Ethernet, 95-98, 159, 581
 ethical hacking, 48, 403. *See also* penetration testing
 ethics (computer), 517-519, 617
 of penetration testing, 49-50
 European Union (EU) Directive on Data Protection,
 501
 Evaluation Assurance Levels. *See* EALs
 evidence, 510-513, 616
 best evidence rule, 511-512
 chain of evidence, 512-513
 credibility of, 510-511
 Fourth Amendment, U.S. Constitution, 513, 616
 hearsay, 511
 imperfect evidence, 512
 proof of authenticity, 511
 exam preparation tips, 621-624
 active study strategies, 622
 common-sense study strategies, 623
 macro and micro study strategies, 622
 pre-testing, 623
 exclusionary rule (Fourth Amendment, U.S.
 Constitution), 513
 execution domains, 351
 Expected Annual Cost (EAC), 408
 expert systems, 261-262
 exposure factor (EF), 200
 external subnets, 132
 extranets, 111
 security management and, 192

F

face recognition, 543. *See also* biometrics
 facility asset class, 534-535
 facsimile transmissions, 147
 fail-closed doors, 543
 fail-open doors, 543
 fail-over clustering, 468, 473-474, 612
 false positives, 189
 Fast Page mode DRAM (FPM DRAM), 257
 FAU (Audit) class, 366
 fault tolerance, 155-156, 474, 584
 fax servers, 432
 FCO (Communication) class, 366
 FCPA (Foreign Corrupt Practices Act), 502
 FCS (Cryptographic Support) class, 366
 FDDI (Fiber Distributed Data Interface), 159
 FDIC (Federal Deposit Insurance Corporation), 444
 FDP (User Data Protection) class, 366
 feasibility studies, 279
 federal laws, 497, 614
 FEMA (Federal Emergency Management Agency), 447,
 451
 FIA (Identification and Authentication) class, 366
 Fiber Distributed Data Interface (FDDI), 159
 fiber-optic cabling, 84-87, 580
 failures, 158
 file system permissions, 389
 file wipe software, 559
 financial computer attacks, 505
 Financial Services Modernization Act, 18
 fingerprinting, 407
 FIPS-197 standard, 314
 fire-related issues
 fire drills, 550
 fire extinguishers, 551-552, 619-620
 fire prevention, 549-550, 619
 fire protection, 549-553

708 INDEX**firewalls, 104-110, 581**

- application-filtering firewalls, 105, 581
- circuit proxy firewalls, 105
- dual-home host firewalls, 109
- dynamic packet filtering firewalls, 106
- kernel proxy firewalls, 107
- packet-filtering firewalls, 104, 107, 581
- screened-host firewalls, 107-108
- screened-subnet firewalls, 108
- stateful packet inspection firewalls, 105-106, 581
- TruSecure testing criteria, 110

flash RAM, 258**floods, 548-549****FMT (Security Management) class, 366****Foreign Corrupt Practices Act (FCPA), 444, 502****foreign keys, 250****forensics, 513-517, 616-617**

- case study, 519-520

forward chaining, 261**Fourth Amendment (U.S. Constitution), 513, 616****FPM DRAM (Fast Page Mode DRAM), 257****FPR (Privacy Class), 366****FPT (Protection of the TSF) class, 366****Frame Relay WAN connections, 116, 159, 583****frames, 76-77****FRU (Resource Utilization) class, 367****Fscan, 406****FTA (TOE Access) class, 367****FTC (Federal Trade Commission), 184****FTP (File Transfer Protocol), 126****FTP (Trusted Path/Channels) class, 367****full backups, 156, 474, 585****full-duplex mode, 96-98****full trust relationships, 40****fusion splices, 87****G****gateways, 110, 582**

- email gateways, 72

GIF (Graphics Interchange Format) files, 73**government**

- classification of data, 220
- regulations, 502
- requirements, security architecture and models, 340-341

Gramm-Leach-Bliley Financial Services Modernization

- Act, 18, 500-502

graphics, 73**grid computing, 262-263****groups, role-based access control, 26-27****grudge-related computer attacks, 505****Guidance Documents (AGD) class, 368****guidelines (information security policies), 210****H****hackers, 275**

- ethical hacking, 48, 403
- Hackers Beware, 50-51

half-duplex mode, 96-98**Halon problems, 551-552****hard-copy libraries, 555-556****hardware backups, 474, 478, 612****hardware-based analyzers, 399****hardware change control, 226-227, 389****hash functions, 316-317, 595****hash totals, 145****hashing, 145****HDLC (High-Level Data-Link Control), 115, 583****HDSL (High-rate Digital Subscriber Line) connections, 117**

- Health Insurance Portability and Accountability Act of 1996 (HIPAA), 18, 183, 500
- hearsay rule, 511, 616
- heat detectors, 560
- HIDSs (host-based intrusion detection systems), 45, 140, 401-403
- hierarchical databases, 252, 590
- Hierarchical Storage Management (HSM), 158, 480
- HIPPA (Health Insurance Portability and Accountability Act of 1996), 18, 183, 500
- host-based intrusion detection systems (HIDSs), 45, 140, 401-403
- hosts, 75
- hot sites, 478, 612
- HSM (Hierarchical Storage Management), 158, 480
- HSSI (High Speed Serial Interface) connections, 118, 583
- hubs, 77, 99-100, 581
- humidity, 548
- Hutt, Arthur, 455
- hybrid sites, 479, 613
- I**
- ICMP (Internet Control Message Protocol), 129 (ICS)²
- Code of Ethics, 518-519
 - physical security categories, 532
 - Web site, 687-688
- identification, 184-185, 577, 585-586
- biometrics, 36
 - case study, 52-53
 - one-time passwords, 36
 - passwords, 35-36
 - single sign-on (SSO) scheme, 37
 - ticket schemes, 36-37
- Identification and Authentication (FIA) class, 366
- IDSs (intrusion detection systems), 44, 401-403
- anomaly detection, 47
 - behavior-based intrusion detection system), 141
 - HIDSs (host-based intrusion detection systems), 45, 140, 401-403
 - knowledge-based intrusion detection system), 140-141
 - NIDSs (network-based intrusion detection systems), 45, 139-141
 - pattern matching, 46-47, 578
 - thresholds, setting, 189
- IEEE 802 standards, 95
- IEEE 802.2 protocol, 76
 - IEEE 802.3 protocol, 76
- IKE (Internet Key Exchange), 372
- illogical processing, 195
- imperfect evidence, 512
- inbound NAT, 143
- incident-response procedures, 211
- incremental backups, 157, 475-476, 585
- indicators, 392-394
- tip-off indicators, 394
- information security management. *See* security management
- information security policies, 205-209, 587
- baselines, creating, 210
 - data access, defining, 209
 - defining, 207
 - development of, 206
 - guidelines, creating, 210
 - inventory of assets, identifying, 207-209
 - procedures, implementing, 210-212
 - standards, setting, 209-210
- information system security standards. *See* security standards
- Information Technology Security Evaluation Criteria. *See* ITSEC

710 INDEX

InfraGard, 507

Initial Program Load (IPL), 397

input controls, 389-390

insurance coverage, 610-611

- planning for insurance claim processing, 461-463
- reviewing in business continuity planning, 460-461, 466

integrity, 182, 311, 585

- access control lists, 347
- Biba security model, 32-33, 345
- Clark-Wilson security model, 346
- networks and, 131

intellectual property laws, 498-499, 613

- copyrights, 498-499
- patents, 498
- trade secrets, 499

inter-repeater links (IRLs), 81

Interagency Contingency Planning Regulation, 444, 610

interference, 182

internal subnets, 132

Internet, 111. *See also* WANs

- Internet-based threats, 416-417, 608
- pseudoflaw attacks, 272
- Web services, 263-265, 591

Internet Control Message Protocol (ICMP), 129

Internet Key Exchange (IKE), 372

Internet Packet Exchange (IPX), 75

Internet Protocol (IP), 75, 129

Internet Protocol Security standard. *See* IPsec

Internet Relay Chat (IRC), 72

Internet Storm Center, 507

interruption vulnerabilities, 535-536

intranets, 111

Introduction and General Model (Common Criteria, Part 1), 364-365

intrusion detection, 43-44, 182, 399-403, 578, 605

- attack signature recognition, 403
- hardware-based analyzers, 399

host versus network intrusions, 44-45

known versus unknown intrusions, 45

packet analysis, 399-400

passive versus active intrusions, 45

physical detection mechanisms, 559-561

resources, 402

software-based sniffers, 399

intrusion detection systems (IDSs), 44, 389, 401-403

anomaly detection, 47

behavior-based IDSs, 141

HIDSs (host-based IDSs), 45, 140, 401-403

NIDSs (network-based IDSs), 45, 139-141

pattern matching, 46-47, 578

thresholds, setting, 189

intrusion prevention, 46

intrusion response, 141-142

inventory of assets, 207-209

ionization-type smoke detectors, 550-551

IP (Internet Protocol), 75, 129

IP addresses, 129

ARP (Address Resolution Protocol), 130

NAT (Network Address Translation), 142-144

RARP (Reverse Address Resolution Protocol), 130

IPL (Initial Program Load), 397

IPsec (IP Security) standard, 124, 135-136, 370-373, 417, 603

architectural components of, 372-373

uses for, 371-372

IPX (Internet Packet Exchange), 75

IRC (Internet Relay Chat), 72

IRL (inter-repeater links), 81

ISDN (Integrated Services Data Network) protocol, 116-117, 583

ISO (International Standards Organization), 69. *See also*

OSI (Open Systems Interconnection) model

isolation

process isolation, 351

resource isolation, 352

ISPTO (U.S. Patent and Trademark Office), 498
 IT roles and responsibilities, 214
 ITSEC (Information Technology Security Evaluation
 Criteria), 360-362, 600-601
 levels, 361-362
 versus the Orange Book standard, 361

J-K

Jack the Ripper (password cracker), 267
 Jammer, 389
 Java applets, 246
 job descriptions, 225
 job rotation, 225
 JPEG (Joint Photographic Experts Group) files, 73
 jump-point attacks, 412

Kerberos, 36-37
 kernel, 388
 kernel proxy firewalls, 107
 security, 284, 352
 key fobs, one-time passwords, 35
 keyboard attacks, 426
 keys (cryptographic), 181, 187, 217-218
 asymmetric encryption, 315-316
 length of, 317-318
 PKI (public key infrastructure), 318-319
 symmetric encryption, 313-314
 keystroke monitoring, 189-190
 keystroke recording, 150
 knowledge engineering, 261
 knowledge-based intrusion detection systems, 140-141
 knowledge-based systems, 261-262
 known attacks, 45
 known-plaintext attacks (KPA), 321-322, 596

L

L1 cache, 257
 L2 cache, 257
 L2TP (Layer 2 Tunneling Protocol), 124
 labeled security protection (Orange Book, class B1), 358
 labels versus access control lists, 353
 LAND attacks, 151
 LANs (local area networks). *See also* network computing
 bridges, 100-101, 581
 data transmission techniques, 94, 580-581
 firewalls, 104-110, 581
 gateways, 110, 582
 hubs, 99-100, 581
 proxies, 110, 582
 repeaters, 99-100, 581
 routers, 103-104, 581
 switches, 100-101, 581
 VLANs (virtual LANs), 101-103, 581
 LAPB (Link Access Procedure Balanced) protocol, 116
 laptops, backups, 476
 lattice of rights, 216
 lattice-based access control, 22-25, 576
 Liptner's lattice, 33, 577
 laws, 496-497
 administrative laws, 497, 613
 civil laws, 497, 613
 criminal laws, 497, 503-505, 613-615
 government regulations, 502
 intellectual property laws, 498-499, 613
 privacy laws, 500-502, 614
 reasonable doubt, 497
 sale and licensing, 499, 613-614
 Layer 2 Tunneling Protocol (L2TP), 124
 layer-3 switches, 75
 layering, 216, 284, 351, 587
 LC4 (password cracker), 267

712 INDEX

LDP (Line Printer Daemon) protocol, 126
 leased lines, 159
 least privilege, 284, 295, 351, 389, 598
 legal evidence, 510-513

- best evidence rule, 511-512
- chain of evidence, 512-513
- credibility of, 510-511
- Fourth Amendment, U.S. Constitution, 513
- hearsay, 511
- imperfect evidence, 512
- proof of authenticity, 511

 Level 1 cache, 257
 Level 2 cache, 257
 licensing, 499, 613-614
 Life Cycle Support (ALC) class, 368
 limited access security mode, 353, 599
 Line Printer Daemon (LPD) protocol, 126
 linear bus topology, 89-91, 580
 Link Access Procedure Balanced (LAPB) protocol, 116
 Liptner's lattice, 33, 577
 LLC sublayer, 76
 load balancing, 156
 lobe ports, 98
 local area networks. *See* LANs
 local hosts, 75
 locks, 542-543
 logging

- access control and, 29
- accountability and, 19

 logic bombs, 243, 271
 logins, limiting attempts, 186
 logs

- access logs, 541-542
- application logs, 398
- audit logs, 542
- security logs, 398
- system logs, 397-398

 loss potential, 196
 losses (risk category), 193

M

MAC (mandatory access control), 21-22, 576
 MAC (Media Access Control) sublayer, 76
 MACs (message authentication codes), 316, 595
 Magic Lantern, 189
 magnetic locking mechanisms, 542
 magnetic sensors, 560
 magnetic tape, 553-555

- degaussing, 558

 Maintenance of Assurance (AMA) class, 368
 malfunctions (risk factor), 193
 malicious code, 274-277, 592
 malware, 243, 246-248
 man-in-the-middle attacks, 323-324, 596
 management. *See also* security management

- controls, 389, 604
- responsibility of, 213

 mandatory access control (MAC), 21-22, 576
 mandatory protection (Orange Book, division B), 358
 masquerading, 148
 massively distributed systems, 245
 maximum tolerable downtime (MTD), 453-454

- validating, 456-457

 MD5 (Message Digest 5) algorithm, 317
 mechanical splices, 87
 meet-in-the-middle attacks, 324, 596
 memory, 256-259
 mesh topology, 93-94, 580
 message authentication codes (MACs), 316, 595
 middleware, 253
 military and intelligence computer attacks, 505
 minimal protection (Orange Book, division D), 358
 mirrored systems, 473-474
 MLS (multilevel secure) security mode, 353, 599
Mobile Agents and Security (Vigna), 248
 mobile sites, 479, 613

modems, 119, 583
 war dialers, 406
modes, security, 352-353, 598-599
monitoring
 audit logs, 395-398
 intrusion detection, 43-44, 399-403, 578
 host versus network intrusions, 44-45
 known versus unknown intrusions, 45
 passive versus active intrusions, 45
 network monitoring, 137-139
 penetration testing, 48, 403-407, 578
 ethical issues, 49-50
 performing, 50-51
 tools, 51
 versus security assessments, 49
monolithic networking model, 68
motion detectors, 560
movie formats, 73
MP3 files, 73
MPDRAM (multiport dynamic RAM), 259
MPEG (Moving Picture Experts Group) files, 73
MTD (maximum tolerable downtime), 453-454
 validating, 456-457
multi-mode fiber-optic cabling, 86
multi-port repeaters, 99. See also hubs
multicasts, 94, 581
multilevel secure (MLS) security mode, 353, 599
multiplexors, 119, 583
mutual aide agreements, 479
MUX (multiplexors), 119

N

NAK (negative acknowledgement) attacks, 272
named perils, 461, 611
NAT (Network Address Translation), 142-144
natural disasters, 446, 539
NCP (Network Core Protocol), 73
NCUA (National Credit Union Administration), 444
negative acknowledgement (NAK) attacks, 272
Nessus vulnerability scanner, 51
Netbus, 273
netcat, 273, 407
NetScan Tools Pro 2001, 406
Network Address Translation (NAT), 142-144
network computing. See also LANs; network security
 10BASE-2 networks, 79-81, 579
 10BASE-5 networks, 79-82, 579
 ARCnet (Attached Resource Computer Network), 99
 cabling
 coaxial, 79-82, 579
 failures, 158
 fiber-optic, 84-87, 580
 UTP (unshielded twisted pair), 82-84, 579-580
 Ethernet, 95-98, 159, 581
 fault tolerance, 155-156, 584
 monolithic networking model, 68
 network databases, 252, 590
 network monitors, 137-139, 399
 network software, 273-274
 OSI model, 68-72, 77-78, 578-579
 Application layer, 72, 126-127, 136-137, 578
 benefits of, 70-71
 Data Link layer, 75-77, 579
 Network layer, 75, 135-136, 579
 Physical layer, 76-77, 579
 Presentation layer, 73, 578
 Session layer, 73-74, 578
 Transport layer, 74, 127-129, 136, 578
 remote access, 119-124, 583-584
 authentication, 124
 dial-up access, 119-120
 tunneling, 120-121
 VPNs (virtual private networks), 121-124
 security management and, 192

714 INDEX

- Token-Ring networks, 98-99
- topologies, 89
 - failures, 159
 - linear bus topology, 89-91, 580
 - mesh topology, 93-94, 580
 - ring topology, 92, 580-581
 - star topology, 91-92, 580
 - tree topology, 93, 580
- VLANs (virtual LANs), 101-103
- WANs (wide area networks), 110-111
 - cell-switched connections, 114, 582
 - circuit-switched connections, 113, 582
 - dedicated connections, 111-112
 - Frame Relay connections, 116, 159
 - HDLC (High-Level Data-Link Control), 115, 583
 - HSSI (High Speed Serial Interface) connections, 118, 583
 - ISDN (Integrated Services Data Network) connections, 116-117
 - packet-switched connections, 113, 582
 - PPP (Point-to-Point Protocol), 114-115
 - SDLC (Synchronous Data-Link Control) connections, 116
 - SMDS (Switched Multimegabit Data Service), 118, 583
 - X.25 connections, 115, 583
 - xDSL (Digital Subscriber Line) connections, 117-118, 583
 - wireless networks, 87-88, 580
- Network Core Protocol (NCP), 73**
- Network File System (NFS) protocol, 74, 126**
- network interface cards, 399**
- network layer, OSI model, 75, 579**
 - protocols, 135-136
- Network Monitor, 399-400**
- network security. *See also* network computing**
 - attacks, 584
 - class A network abuses, 147-148
 - class B network abuses, 148-149
 - class C network abuses, 149-150
 - class D network abuses, 150-151
 - class E network abuses, 152-154
 - class F network abuses, 154-155
 - case study, 160-162
 - common attributes, 79
 - email, 146-147
 - exam objective overview, 61-67
 - firewalls, 104-110, 581
 - hashing, 145
 - intrusion detection, 139-141
 - intrusion response, 141-142
 - NAT (Network Address Translation), 142-144
 - network monitoring, 137-139
 - port-based, 103
 - printers, 147
 - SANs (storage area networks), 259-260, 591
 - security boundaries, 132-133
 - transparency, 144-145
 - tunneling as, 120-121
 - wireless networks, 88
- network services clustering, 156, 585**
- network sniffing, 149**
- network-based IDSs (intrusion detection systems), 45, 401-403**
- newsgroups, 72**
- NFS (Network File System) protocol, 74, 126**
- NICs (network interface cards), 79**
- Nimda, 246-247**
- NMAP, 406**
- nmap port scanner, 52**
- NNTP (Network News Transfer Protocol), 72**
- noninference models (access control), 33, 577**

nondistributed systems, 241-243
 nonrepudiation, 188, 312
 nonvolatile storage, 256
 nuke attacks, 412
 Nuke Nabber, 402

O

object-oriented databases, 252, 590
 object-oriented programming, 289-291, 594
 abstraction, 217
 objects (Bell-LaPadula security model), 343-344
 OC-x (Optical Carrier X) connections, 112
 OCC (Office of the Comptroller), 444
 offsite data storage, 559
 one-time ciphers, 318, 596
 one-time passwords, 35-36
 one-way trust relationships, 40
 open systems, 350-351
 operating systems
 fingerprinting, 407
 privileges associated with, 388
 operational controls, 389
 operations security, 603-604
 administrative management, 418-420, 608
 best practices, 420-428
 antiviral controls, 423-425
 change management control, 427-428
 privileged operation functions, 421-422
 protecting sensitive information, 425-427
 countermeasures to threats, 408
 disgruntled employees, 415-416
 employee-related threats, 412-414
 hiring and firing/exit practices, 414-415
 information system threats, 409-410
 Internet-based threats, 416-417
 mainframe threats, 410-411
 physical threats, 417-418
 threat risk analysis, 408-409
 monitoring, 395-398
 audit logs, 395-398
 intrusion detection, 399-403
 penetration testing, 403-407
 objective overview, 381-382
 overview of, 385-386
 process of, 392-395
 roles of, 387-395
 identifying available controls, 389-391
 identifying privileges to be restricted, 388
 identifying resources to be protected, 387-388
 suggested readings, 438
 OPSEC. *See operations security*
 Orange Book standard, 133-134, 356, 599-600. *See also*
 Rainbow Series
 classifications, 357-358
 criticisms of, 358-359
 versus ITSEC, 361
 OSI (Open Systems Interconnection) model, 68-72,
 77-78, 578-579
 Application layer, 72, 126-127, 136-137, 578
 benefits of, 70-71
 Data Link layer, 75-77, 579
 Network layer, 75, 135-136, 579
 Physical layer, 76-77, 579
 Presentation layer, 73, 578
 Session layer, 73-74, 578
 Transport layer, 74, 127-129, 136, 578
 OTS (Office of Thrift Supervision), 444
 output controls, 389-390

716 INDEX

P

- packed databases, 557
- packet analysis, 399-401
- packet-filtering firewalls, 104, 107, 581
- packet-sniffing software, 137-139
- packet-switched connections, WANs, 113, 582
- paper documents, storage of, 555-556
- parity checks, 316
- partial backups, 474
- passive
 - attacks, 45
 - monitoring, 189
 - physical access controls, 542-544
 - technologies, 90
- passwords, 35-36, 185-187
 - access control administration, 28-29
 - checkers, 186
 - brute-force attacks, 148, 266-267
 - dictionary attacks, 266-267
 - generators, 186
 - one-time passwords, 35-36
- patents, 498, 613
- pattern-based application recognition, 138
- pattern-matching IDSs (intrusion detection systems), 46-47, 578
- PBX fraud, 148-149
- PBX Scanner, 406
- PCMCIA cards, 258
- PDA (personal digital assistants), backups, 476
- PEM (Privacy Enhanced Mail) protocol, 137
- penetration testing, 48, 403-407, 578, 605
 - ethical issues, 49-50
 - performing, 50-51
 - tools for, 51-52
 - versus security assessments, 49
- permissions
 - access control and, 29
 - permission sets, 388
- PGP (Pretty Good Privacy), 218
- phone tag (war dialer), 406
- photoelectric smoke detectors, 550-551
- phreaking, 275
- physical and components asset class, 534-535
- physical and environmental procedures, 211
- physical intrusion-detection mechanisms, 559-561
- Physical layer, OSI model, 76-77, 579
- physical security, 532-533
 - asset classes, 533-535, 617-618
 - case study, 562-563
 - detection mechanisms, 559-561
 - exam objective overview, 529-532
 - paper storage, 555-556
 - removable electronic media, 553-555
 - secure sites, 538
 - environmental controls, 547-548, 619
 - fire prevention and detection, 549-553, 619
 - location and construction of sites, 539-540, 561-562
 - location and construction of sites, 618
 - physical access controls, 540-544, 618
 - power-supply issues, 544-547, 618-619
 - water exposure problems, 548-549, 619
 - theft, 537, 618
 - vulnerabilities, 535-537
- physical threats, 417-418
- piggy-backing, 154
- PIN codes, 185
- PKI (public key infrastructure), 187, 318-319, 596
- plaintext, 313, 595
 - chosen-plaintext attacks, 322, 596
 - known-plaintext attacks, 321-322, 596
- planning for information security, 191-192

- Point-to-Point Tunneling Protocol (PPTP), 123
- policies (information security), 205-209
 - baselines, creating, 210
 - data access, defining, 209
 - defining, 207
 - development of, 206
 - guidelines, creating, 210
 - inventory of assets, identifying, 207-209
 - procedures, implementing, 210-212
 - standards, setting, 209-210
- policy routing, 143
- POP3 (Post Office Protocol 3), 126
- port-based security, 103
- port mirroring, 139
- port scanning, 51, 139, 154, 406
- positive air pressure, 426
- Post Office Protocol 3 (POP3), 126
- power supplies, site security and, 544-547, 618-619
- PP (Protection Profile), 365
- PPP (Point-to-Point Protocol), 114-115, 583
- PPTP (Point-to-Point Tunneling Protocol), 123
- Presentation layer, OSI model, 73, 578
- Pretty Good Privacy (PGP), 218
- preventative controls, 390, 604
- PRI (primary rate interface), 117
- primary keys, 250
- primary storage, 256
- principle of least privilege, 17, 295
- principles, security, 351-352
- printer security, 147
- privacy, 183-184, 585
 - Privacy Act of 1974, 183
 - privacy law, 500-502, 614
- Privacy (FPR) class, 366
- Privacy Enhanced Mail (PEM) protocol, 137
- private data, 219, 588
- private key encryption, 217-218
- privileged instruction, 388
- privileged operations, 421-422
- privileges, restricting, 388
- probing attacks, 154-155
- procedures (information security policies), 210-212
- processes
 - isolation, 351, 598
 - layering, 351
 - least privilege, 351
- promiscuous mode, 399
- proof beyond a reasonable doubt, 497
- proof by a preponderance of the evidence, 497
- proprietary data, 219, 588
- protection mechanisms, 215-218, 587
 - abstraction, 217
 - data hiding, 217
 - encryption, 217-218
 - layering, 216
- Protection of the TSF (FPT) class, 366
- Protection Profile (PP), 365
- Protection Profile Evaluation (PPE) class, 367
- protocol filters, 399
- protocols, networking, 125-130
- prototyping, 594
- proving the chain of custody, 190
- proxies, 110, 582
- pseudoflaw attacks, 272
- public data, 219, 588
- public key cryptography, 187, 217-218
- public key encryption, 315-316
- public key infrastructure (PKI), 187, 318-319, 596
- public requirements, security architecture and models, 340-341

718 INDEX

Q

QADAD (quick-and-dirty application development), 283

QIC (Quarter Inch Cartridge) backup systems, 157

qualitative risk analysis, 196-197, 202, 409, 586

quantitative risk analysis, 196-197, 408

QuickTime, 73

R

RAD (Rapid Application Development), 282-283, 593

radiation technologies, 561

RADIUS (Remote Authentication Dial-In User Service), 37-38, 124, 577, 584

RAID (redundant array of inexpensive disks), 155-156, 474, 584-585

RAID 0, 155

RAID 1, 155, 585

RAID 2, 155, 585

RAID 3, 155, 585

RAID 4, 155, 585

RAID 5, 156, 585

Rainbow Series, 133, 359-360. *See also* Orange Book standard

RAM (random access memory), 256-259

removing data from, 425

RAMBUS DRAM, 257

random attacks, 412

Rapid Application Development (RAD), 282-283

RARP (Reverse Address Resolution Protocol), 130

real evidence, 510

real memory, 256

Real Secure Desktop Protection, 402

Real Secure Server Sensor, 402

real-time clock (RTC), 258

RealAudio files, 73

RealSecure for Nokia, 402

reasonable doubt, 497

recovery time objective (RTO), 453-454

Red Book (Rainbow Series), 359

redundant array of inexpensive disks. *See* RAID

redundant sites, 479, 613

reference monitor, 348-349

referential integrity, 250

reflexive property (lattice-based access control), 23-25

registers, 257

relational databases, 251, 590

remote access, 119-124, 583-584

authentication, 124, 577

tunneling, 120-121

VPNs (virtual private networks), 121-124

Remote Authentication Dial-In User Service. *See* RADIUS

remote-control lock systems, 542

remote hosts, 75

Remote Procedure Calls (RPCs), 74

removable electronic media, 553-555

removal vulnerabilities, 535-536

repeaters, 77, 99-100, 581

replacement cost, 461

replay attacks, 323, 596

residual risk, 204

Resource Utilization (FRU) class, 367

resources

identifying resources to be protected, 387-388

isolation, 352

listing necessary resources, 469-470

Reverse Address Resolution Protocol (RARP), 130

revision control, 226-227

RFC 1087, "Ethics and the Internet," 517-518

RI/RO (Ring in/Ring out) ports, 98

Rijndael algorithm, 314

ring 0, 388

ring topology, 92, 580-581

risk analysis, 194-195, 408-409, 586

- asset valuation, 196-197
- cost-benefit analysis, 194-195, 203-204
- government vs. nongovernment organizations, 194
- outside consultants, 197
- qualitative risk analysis, 202, 586
- quantitative vs. qualitative approaches, 196-197
- responses to, 587
- steps of, 197-200
- threats and vulnerabilities, identifying, 195-196
- variables, 200

risk management, 192-205, 586

- countermeasures, cost/benefit analysis, 203-204
- risk analysis, 194-195
 - asset valuation, 196-197
 - cost-benefit analysis, 194-195, 203-204
 - government vs. nongovernment organizations, 194
 - outside consultants, 197
 - qualitative risk analysis, 202
 - quantitative vs. qualitative approaches, 196-197
 - steps of, 197-200
 - threats and vulnerabilities, identifying, 195-196
 - variables, 200
- risk categories, 193
- risk factors, 193

role-based access control, 26-27, 576

roles

- IT roles and responsibilities, 214
- operations security roles, 387-395
 - identifying available controls, 389-391
 - identifying privileges to be restricted, 388
 - identifying resources to be protected, 387-388
- security roles and responsibilities, 212
 - integration of, 214-215
 - IT staff, 214
 - management responsibilities, 213
 - user responsibilities, 213-214

ROM (read-only memory), removing data from, 425

routers, 75, 103-104, 118, 581-583

routing, 75

RPCs (remote procedure calls), 74

RTC (real-time clock), 258

RTO (recovery time objective), 453-454

rule-based access control, 25, 576. *See also* ACLs (access control lists)

S

S/Key one-time password program, 36

S/MIME (Secure/Multipurpose Internet Mail Extensions) protocol, 137

sabotage, 446

sadadmin worm, 247

safe harbor, 501, 614

Samspade, 407

Sandtap, 406

SANs (storage area networks), 259-260, 591

SANS Institute, 507

SBU (sensitive but unclassified) data, 220, 588

SC (Stick and Click) connectors, 86

scope, determining

- of business continuity plans, 451-452
- of disaster recovery plans, 468

screened subnets, 132

screened-host firewalls, 107-108

screened-subnet firewalls, 108

SDLC (Synchronous Data-Link Control) protocol, 116, 583

SDRAM (synchronous DRAM), 257

SDSL (Single-line Digital Subscriber Line) connections, 117

secondary storage, 256

secret data, 220, 588

secret-key encryption, 313-314

720 INDEX

- Secure Electronic Transmission (SET)**, 137
- secure sites.** *See* site security
- Secure Socket Layer (SSL)**, 136, 218
- SecureID, one-time passwords**, 35
- security architecture**, 283-285, 348, 597-599
 - covert channels, 354, 599
 - labels versus access control lists, 353
 - modes, 352-353, 598-599
 - open versus closed systems, 350-351
 - reference monitor, 348-349
 - requirements, government versus public, 340-341
 - security principles, 351-352, 598
- security assessments**, 578
 - versus penetration testing, 49
- Security Assurance Requirements (Common Criteria, Part 3)**, 367-368
- security awareness training**, 227-228
- security domains (Orange Book, class B3)**, 358
- Security Functional Requirements (Common Criteria, Part 2)**, 365-367
- security incidents**, 505-506
 - advance planning, 506-507
 - investigation of, 507-509
 - legal evidence, 510-513
 - best evidence rule, 511-512
 - chain of evidence, 512-513
 - credibility of, 510-511
 - hearsay, 511
 - proof of authenticity, 511
 - U.S. Fourth Amendment, 513
 - major categories, 505-506
- security kernels**, 284, 352
- security logs**, 398
- security management**, 179-180. *See also* security roles and responsibilities
 - change control, 226-227
 - classifying data, 218-222
 - commercial classification, 219
 - criteria for, 221
 - government classification, 220
 - procedures, 221-222
 - employment policies, 222-225
 - information security policies, 205-209
 - baselines, creating, 210
 - data access, defining, 209
 - defining, 207
 - development of, 206
 - guidelines, creating, 210
 - inventory of assets, identifying, 207-209
 - procedures, implementing, 210-212
 - standards, setting, 209-210
 - planning for information security, 191-192
 - principles, 180-183
 - accountability, 188
 - auditing, 188-190
 - authentication, 184-187
 - availability, 183
 - confidentiality, 181-182
 - documentation, 190-191
 - identification, 184-185
 - integrity, 182
 - nonrepudiation, 188
 - privacy, 183-184
 - protection mechanisms, 215-218
 - abstraction, 217
 - data hiding, 217
 - encryption, 217-218
 - layering, 216
 - risk management, 192-205
 - countermeasures, 203-204
 - risk analysis. *See* risk analysis
 - risk categories, 193
 - risk factors, 193
 - security awareness training, 227-228
 - separation of duties, 390
- Security Management (FMT) class**, 366
- security managers**, 419-420

- security models, 342**
 - access control lists, 347, 597
 - Bell-LaPadula, 30-33, 216, 343-345, 576
 - Biba, 32-33, 345, 577
 - Clark-Wilson, 346, 597
 - summary of, 347-348
- security modes, 352-353**
- security of operations, 603-604**
 - administrative management, 418-420, 608
 - best practices, 420-428
 - antiviral controls, 423-425
 - change management control, 427-428
 - privileged operation functions, 421-422
 - protecting sensitive information, 425-427
 - countermeasures to threats, 408
 - disgruntled employees, 415-416
 - employee-related threats, 412-414
 - hiring and firing/exit practices, 414-415
 - information system threats, 409-410
 - Internet-based threats, 416-417
 - mainframe threats, 410-411
 - physical threats, 417-418
 - threat risk analysis, 408-409
 - monitoring, 395-398
 - audit logs, 395-398
 - intrusion detection, 399-403
 - penetration testing, 403-407
 - objective overview, 381-382
 - overview of, 385-386
 - process of, 392-395
 - roles of, 387-395
 - identifying available controls, 389-391
 - identifying privileges to be restricted, 388
 - identifying resources to be protected, 387-388
 - suggested readings, 438
- Security Parameters Index (SPI), 372**
- security perimeter, 352**
- security principles, 351-352, 598**
- Security Reference Monitor (SRM), 349**
- security roles and responsibilities, 212**
 - integration of, 214-215
 - IT staff, 214
 - management responsibilities, 213
 - user responsibilities, 213-214
- security standards, 355-356**
 - CC (Common Criteria), 362-369, 601-603
 - areas not addressed by, 369
 - EALs (Evaluation Assurance Levels), 368-369, 602
 - objectives, 363
 - Part 1 (Introduction and General Model), 364-365
 - Part 2 (Security Functional Requirements), 365-367
 - Part 3 (Security Assurance Requirements), 367-368
 - ITSEC (Information Technology Security Evaluation Criteria), 360-362, 600-601
 - Orange Book. *See* Orange Book standard
 - summary of, 370
 - TCSEC (Trusted Computer Security Evaluation Criteria), 356-360
- Security Target (ST), 365-367**
- Security Target Evaluation (ASE) class, 367**
- segmentation, 101**
 - and reassembly, 74
- segregation of duties, 512**
- semantic integrity, 250**
- sensitive but unclassified (SBU) data, 220, 588**
- sensitive data, 219, 588, 609**
- sensitive information, protecting, 425-427**
- sensors, 559-561**
- separation of duties, 29, 284, 390**
- sequential access, 256**
- Server Message Block (SMB), 73**
- servers**
 - access servers, 119, 583
 - antiviral controls, 424

722 INDEX

- edge servers, 277
- fax servers, 432
- standby servers, 468
- Session layer, OSI model, 73-74, 578
- SET (Secure Electronic Transmission), 137
- setting the clipping level, 403, 605
- shunt trips, 471
- signature-matching intrusion detection systems, 46-47
- Simple Key Management for Internet Protocol (SKIP), 136
- Simple Mail Transfer Protocol (SMTP), 126
- Simple Network Management Protocol (SNMP), 127
- simple security rule
 - Bell-LaPadula security model, 30-31
 - Biba security model, 32
- simplex, 85
- single sign-on (SSO) scheme, 37
- single-key encryption, 313-314
- single-loss expectancy (SLE), 199-201
- single-mode fiber-optic cabling, 86-87
- site security, 538
 - environmental controls, 547-548
 - fire prevention and detection, 549-553
 - location and construction of sites, 539-540, 561-562
 - physical access controls, 540-544
 - power-supply issues, 544-547
 - water exposure problems, 548-549
- site-to-site VPNs, 122
- SKIP (Simple Key Management for Internet Protocol), 136
- SLE (single-loss expectancy), 199-201
- SLIP (Serial Line Internet Protocol), 583
- smart cards, 542-543
- SMB (Server Message Block), 73
- SMBRelay attack, 270
- SMDS (Switched Multimegabit Data Service), 118, 583
- smoke detectors, 550-551
- smoking policies, 550
- SMTP (Simple Mail Transfer Protocol), 126
- SMURF attacks, 151, 268
- sniffers, 274, 399, 560
 - switched networks, 400
- sniffing, 43, 155, 578
 - network sniffing, 149
 - packet-sniffing software, 137-139
- SNMP (Simple Network Management Protocol), 127
- social engineering, 148
 - as attack to confidentiality, 181
- software. *See also* systems development
 - antivirus software, 277
 - attacks. *See* attacks
 - code libraries, 264
 - coding practices, 286-296
 - distributed systems, 244-245
 - malware for, 246-248
 - massively distributed systems, 245
 - file wipe software, 559
 - forensics software, 516
 - illegitimate use of, 272-273
 - network software, 273-274
 - nondistributed systems, 241-243
 - sale and licensing, 499
 - vulnerabilities, 270-272
- software change control, 227
- software-based sniffers, 399
- sound formats, 73
- spam, 146
- SPI (Security Parameters Index), 372
- spikes, 545-546, 618
- spiral system development lifecycle model, 280-282, 592-593
- split tunneling, 122
- spoofing attacks, 42-43, 152, 269-270, 578
- sprinkler systems, 553
- SQL (Structured Query Language), 74
- SRAM (static RAM), 257

- SRM (Security Reference Monitor), 349
- SSL (Secure Socket Layer), 136, 218
- SSO (single sign-on) scheme, 37
- ST (Security Target), 365-367
- ST (Stick and Turn) connectors, 86
- standards (information security policies), 209-210
- standby servers, 468
- star (*) property
 - Bell-LaPadula security model, 31-32, 345
 - Biba security model, 32
- star topology, 91-92, 580
- state laws, 497
- stateful packet inspection firewalls, 105-106, 581
- static electricity, 545-546
- static RAM (SRAM), 257
- station ports, 98
- storage area networks (SANs), 259-260, 591
- storage (data), 256-259
 - document libraries, 555-556
 - electronic media, 553-555
 - offsite, 559
 - RAID (redundant array of inexpensive disks), 155-156, 474, 584-585
 - SANs (storage area networks), 259-260
- strategic attacks, 412
- strong authentication, 185
- strong encryption, 321
- structured programming, 286-289, 594
- structured protection (Orange Book, class B2), 358
- Structured Query Language (SQL), 74
- study tips, CISSP exam, 621-624
 - active strategies, 622
 - common-sense strategies, 623
 - macro and micro strategies, 622
 - pre-testing, 623
- subjects
 - access control list subjects, 347
 - Bell-LaPadula security model, 30, 343-344
 - subnets, 132-133
 - Superscan, 406
 - supplies and materials asset class, 534-535
 - support asset class, 534-536
 - surges, 545-546, 618
 - Sustainable Computing Consortium, 239
 - SWIPE protocol, 136
 - Switched Multimegabit Data Service (SMDS), 118
 - switched networks, sniffing on, 400
 - switches, 76, 100-101, 581
 - layer-3 switches, 75
 - WAN switches, 118, 583
 - symmetric encryption, 217, 313-314, 595
 - SYN flooding, 150
 - synchronous DRAM (SDRAM), 257
 - synchronous serial connections, 111
 - synchronous tokens, 187
 - SYNs (synchronizations), 127
 - system high security mode, 284, 352, 598
 - system logs, 397-398
 - system low security mode, 284
 - systems development, 239-240
 - case study, 296-297
 - coding practices, 286, 292-296
 - CASE (computer-aided software engineering), 291-292, 594
 - object-oriented programming, 289-291, 594
 - structured programming, 286-289, 594
 - controls, 277-285
 - best practices, 285
 - RAD (Rapid Application Development), 282-283, 593
 - security control architecture, 283-285, 593
 - spiral lifecycle model, 280-282, 592-593
 - waterfall lifecycle model, 278-280, 592
 - databases, 249-255
 - data models, 251-252
 - versus data marts, 255
 - versus data warehouses, 255

724 INDEX

distributed systems, 244-245
 malware for, 246-248
 massively distributed systems, 245
 exam objective overview, 235-239
 knowledge-based systems, 261-262
 malicious code, 274-277
 memory issues, 256-259
 nondistributed systems, 241-243
 software attacks
 antivirus software, 277
 brute-force attacks, 266
 DDoS (distributed denial-of-service) attacks, 269
 dictionary attacks, 266-267
 DoS (denial-of-service) attacks, 267-269
 hidden code, 270
 illegitimate software use, 272-273
 logic bombs, 271
 NAK (negative acknowledgement) attacks, 272
 network software, 273-274
 pseudoflaw attacks, 272
 spoofing, 269-270
 trap doors, 271

T

T1 connections, 112
T3 connections, 112
TACACS (Terminal Access Controller Access Control System), 37-38, 124
TACACS+, 124, 577, 584
tape arrays, 158
tape drives, 553-555
Target of Evaluation (TOE), 361-368
TCB (Trusted Computing Base), 351
TCP (Transport Control Protocol), 128
 TCP hijacking, 153
 TCPWrappers, 402

TCP/IP (Transmission Control Protocol/Internet Protocol), 125, 584
 Application Layer protocols, 126-127
 Internet Layer protocols, 129-130
 Transport Layer protocols, 127-129
TCSEC (Trusted Computer Security Evaluation Criteria), 133-134, 356-360
TDRs (time domain reflectometers), 82, 579
teardrop attacks, 151
technical controls, 389, 604
telecommunications and network security. See network security
Telnet, 127
temporary memory. See RAM
temporary passwords, 28-29
Terminal Access Controller Access Control System. See TACACS
termination of employees, 224-225
terrorist-related computer attacks, 446, 505, 561-562
Tests (ATE) class, 368
TFTP (Trivial File Transfer Protocol), 127
theft, 537, 618
thick net, 79-82
thin net, 79-81
threat agents, 196
threats, 392, 605-606. *See also vulnerabilities*
 countermeasures, 411-412
 disgruntled employees, 415-416
 employee-related threats, 412-414
 hiring and firing/exit practices, 414-415
 Internet-based threats, 416-417
 physical threats, 417-418
 exam objective overview, 529-532
 identifying, 195-196
 information system threats, 409-410
 mainframe threats, 410-411
 risk analysis, 408-409
thresholds (auditing), 189

tickets, 36-37
 TIFF (Tag Image File Format) files, 73
 time domain reflectometers (TDRs), 82, 579
 Time of Check to Time of Use (TOC/TOU), 271-272
 time-modulated ultra-wide band, 561
 tip-off indicators, 394
 TLS (Transport Layer Security) protocol, 136
 TOC/TOU (Time of Check to Time of Use), 271-272
 TOE (Target of Evaluation), 361-368
 TOE Access (FTA) class, 367
 token devices, 187
 token passing, 98, 581
 Token-Ring networks, 98-99, 159
 token systems, 542
 top secret data, 220, 588
 trade secret law, 499, 613
 transitive property (lattice-based access control), 23-25
 transparency, 144-145
 Transport Control Protocol. *See* TCP
 Transport Layer Security (TLS) protocol, 136
 Transport layer, OSI model, 74, 578
 protocols, 127-129, 136
 trap doors, 271
 tree topology, 93, 580
 Triple-DES, 218, 314
 Tripwire, 389, 403
 Trivial File Transfer Protocol (TFTP), 127
 Trojan horses, 152-153, 243, 247
 troubleshooting network failures, 158-159
 trust relationships, 40
 Trusted Computer Security Evaluation Criteria. *See* TCSEC
 Trusted Computing Base (TCB), 351
 Trusted Path/Channels (FTP) class, 367
 trusted subjects (Bell-LaPadula security model), 345
 tunneling, 120-121, 583
 split tunneling, 122

two-key encryption, 315-316
 type enforcement, 132-133

U

U.S. Commerce Department's Bureau of Export Administration (BXA), 502
 U.S. Constitution, Fourth Amendment, 513, 616
 U.S. Export Administration Regulations, 502
 U.S. Federal Financial Examination Council, 444
 U.S. Federal Privacy Act of 1974, 183
 U.S. Federal Trade Commission (FTC), 184
 U.S. Patent and Trademark Office (USPTO), 498
 UDP (User Datagram Protocol), 127-129
 port scanning, 406
 unauthorized access banners, 504-505
 unclassified data, 220, 588
 unicasts, 76, 94, 580
 unknown attacks, 45
 unshielded twisted pair (UTP) cabling, 82-84, 158, 579-580
 untrusted subjects (Bell-LaPadula security model), 345
 UPS (uninterruptible power supply), 543-547
 User Data Protection (UDP) class, 366
 User Datagram Protocol (UDP), 127-129
 user IDs, 28, 34. *See also* authentication; identification
 challenge response schemes, 36
 sniffing, 43
 ticket schemes, 36-37
 users, security responsibilities, 213-214
 user's space, 216
 UTP (unshielded twisted pair) cabling, 82-84, 579-580
 troubleshooting, 84, 158

726 INDEX

V

valuable papers, insurance coverage, 461
 Vampire taps, 81
 variance detection controls, 389, 604
 VDSL (Very High Digital Subscriber Line) connections, 118
 verified design (Orange Book, class A1), 358
 vibration sensors, 560
 video RAM (VRAM), 259
 virtual circuits, 74
 virtual LANs (VLANs), 101-103, 581
 virtual memory, 256, 283
 virtual nodes, 122
 virtual private networks. *See* VPNs
 viruses, 243

- email and, 146-147
- network security and, 153

 vital records, backing up, 477
 VLANs (virtual LANs), 101-103
 volatile memory. *See* RAM
 VPNs (virtual private networks), 121-124, 584

- encryption, 218
- security management and, 192

 VRAM (video RAM), 259
 vulnerabilities, 392, 535-537. *See also* threats

- destruction vulnerabilities, 535-536
- disclosure vulnerabilities, 535-536
- exam objective overview, 529-532
- identifying, 195-196
- interrupting vulnerabilities, 535-536
- removal vulnerabilities, 535-536
- software-based vulnerabilities, 270-272

 Vulnerability Assessment (AVA) class, 368
 vulnerability scanners, 51

W

WAN switches, 118, 583
 WANs (wide area networks), 110-111. *See also* network computing; network security

- cell-switched connections, 114
- circuit-switched connections, 113
- dedicated connections, 111-112
- Frame Relay connections, 116, 583
- HDLC (High-Level Data-Link Control), 115, 583
- ISDN (Integrated Services Digital Network), 583
- packet-switched connections, 113
- PPP (Point-to-Point Protocol), 114-115
- SDLC (Synchronous Data-Link Control), 116, 583
- SMDS (Switched Multimegabit Data Services), 583
- X.25 connections, 115

 war, 446
 war dialers, 406
 warm sites, 478, 613
 waste disposal, 556-559, 620
 water exposure problems, 548-549, 619
 waterfall system development lifecycle model, 278-280, 592
 WAV (Windows Audio Volume) files, 73
 Web services, 263-265, 591
 wet standpipe systems, 553
 wide area networks. *See* WANs
 windowing, 128
 windows, 544
 wireless networks, 87-88, 580
 Wiretap Act, 504, 615
 WMF (Windows Media File) format, 73
 work for hire, 499, 613
 worms, 246

- network security and, 153

 Writing Secure Code (Howard and LeBlanc), 238, 293
 WWW (World Wide Web) applications, 72

X-Z

X Window protocol, 127

X-rays, 561

X.25 WAN connections, 115, 583

xDSL (Digital Subscriber Line) connections, 117-118,
583

XML (Extensible Markup Language), 264


Zip drives, 157

zipcords, 85

zombies, 151

SANS
P R E S S

**MASTER THE TOOLS OF THE
 NETWORK SECURITY TRADE WITH THE
 OFFICIAL BOOK FROM SANS PRESS!**



**SANS GIAC Certification:
 Security Essentials Toolkit (GSEC)**

Eric Cole, Mathew Newfield, John M. Millican
 with foreword by Stephen Northcutt
 0-7357-2774-9 • 384 pages • \$49.99 US

QUE[®]
 CERTIFICATION

You need more than a hammer to build a house, and you need more than one tool to secure your network. *Security Essentials Toolkit* covers the critical tools that you need to secure your site, showing you why, when, and how to use them. Based on the SANS Institute's renowned Global Information Assurance Certification (GIAC) program, this book takes a workbook-style approach that gives you hands-on experience and teaches you how to install, configure, and run the best security tools of the trade.

www.quepublishing.com

informIT

www.informit.com

Your Guide to Information Technology Training and Reference

Que has partnered with **InformIT.com** to bring technical information to your desktop. Drawing on Que authors and reviewers to provide additional information on topics you're interested in, **InformIT.com** has free, in-depth information you won't find anywhere else.

Articles

Keep your edge with thousands of free articles, in-depth features, interviews, and information technology reference recommendations – all written by experts you know and trust.

Online Books

Answers in an instant from **InformIT Online Books'**
600+ fully searchable online books. Sign up now and
get your first 14 days **free**.

POWERED BY

Safari

Catalog

Review online sample chapters and author biographies to choose exactly the right book from a selection of more than 5,000 titles.

As an **InformIT** partner, **Que** has shared the knowledge and hands-on advice of our authors with you online. Visit **InformIT.com** to see what you are missing.

QUE[®]

www.quepublishing.com

What if Que
 joined forces to deliver the
 best technology books in a
 common digital reference platform?

We have. Introducing
InformIT Online Books
 powered by Safari.

■ **Specific answers to specific questions.**

InformIT Online Books' powerful search engine gives you
 relevance-ranked results in a matter of seconds.

■ **Immediate results.**

With InformIT Online Books, you can select the book you
 want and view the chapter or section you need immediately.

■ **Cut, paste, and annotate.**

Paste code to save time and eliminate typographical errors.
 Make notes on the material you find useful and choose
 whether or not to share them with your workgroup.

■ **Customized for your enterprise.**

Customize a library for you, your department, or your entire
 organization. You pay only for what you need.

Get your first 14 days FREE!

InformIT Online Books is offering its members a 10-book subscription risk free
 for 14 days. Visit <http://www.informit.com/onlinebooks> for details.


POWERED BY
safari
informIT
Online Books

As an InformIT partner,
 Que has shared the knowl-
 edge and hands-on advice
 of our authors with you
 online. Visit InformIT.com to
 see what you are missing.

informit.com/onlinebooks

Test. Get Certified!

You have the experience and the training — now demonstrate your expertise and get the recognition your skills deserve. An IT certification increases your credibility in the marketplace and is tangible evidence that you have the know-how to provide top-notch support to your employer.



Why Test with VUE?

Using the speed and reliability of the Internet, the most advanced technology and our commitment to unparalleled service, VUE provides a quick, flexible way to meet your testing needs.

Three easy ways to register for your next exam, all in real time:

- ◆ Register online at www.vue.com
- ◆ Contact your local VUE testing center. There are over 3000 quality VUE testing centers in more than 130 countries. Visit www.vue.com for the location of a center near you.
- ◆ Call a VUE call center. In North America, call toll-free 800-TEST-NOW (800-837-8734). For a complete listing of worldwide call center telephone numbers, visit www.vue.com.

Call your local VUE testing center and ask about TESTNOW!™ same-day exam registration!

The VUE testing system is built with the best technology and backed by even better service. Your exam will be ready when you expect it and your results will be quickly and accurately transmitted to the testing sponsor. Test with confidence!

Get
Certified!

Visit www.vue.com
for a complete listing of
IT certification exams
offered by VUE



When IT really matters...Test with VUE!

CramSession

– the difference between Pass
... or Fail

CramSession.com is #1 for IT Certification on the Net.

There's no better way to prepare for success in the IT industry. Find the best IT certification study materials and technical information at CramSession. Find a community of hundreds of thousands of IT pros just like you who help each other pass exams, solve real-world problems, and discover friends and peers across the globe.

CramSession – #1 Rated Certification Site!

- #1 by *TechRepublic.com*
- #1 by *TechTarget.com*
- #1 by *CertMag's Guide to Web Resources*



Visit Cramsession.com today!

...and take advantage of the best
IT learning resources.



CramSession has IT all!

- **The #1 study guides on the Net.** With over 250 study guides for IT certification exams, we are the Web site every techie visits before passing an IT certification exam.
- **Practice questions.** Get the answers and explanations with our CramChallenge practice questions delivered to you daily.
- **The most popular IT forums.** CramSession has over 400 discussion boards loaded with certification information where our subscribers study hard, work hard, and play harder.
- **e-Newsletters.** Our IT e-Newsletters are written by techs for techs: IT certification, technology, humor, career, and more.
- **Technical papers and product reviews.** Find thousands of technical articles and whitepapers written by industry leaders, trainers, and IT veterans.
- **Exam reviews.** Get the inside scoop before you take that expensive certification exam.
- **And so much more!**

"On top of everything else, I find the best deals on training products and services for our CramSession members,"

Jami Costin,
Product Specialist



CramSession
Prepare for Success!

www.cramsession.com